Robotics Research Technical Report

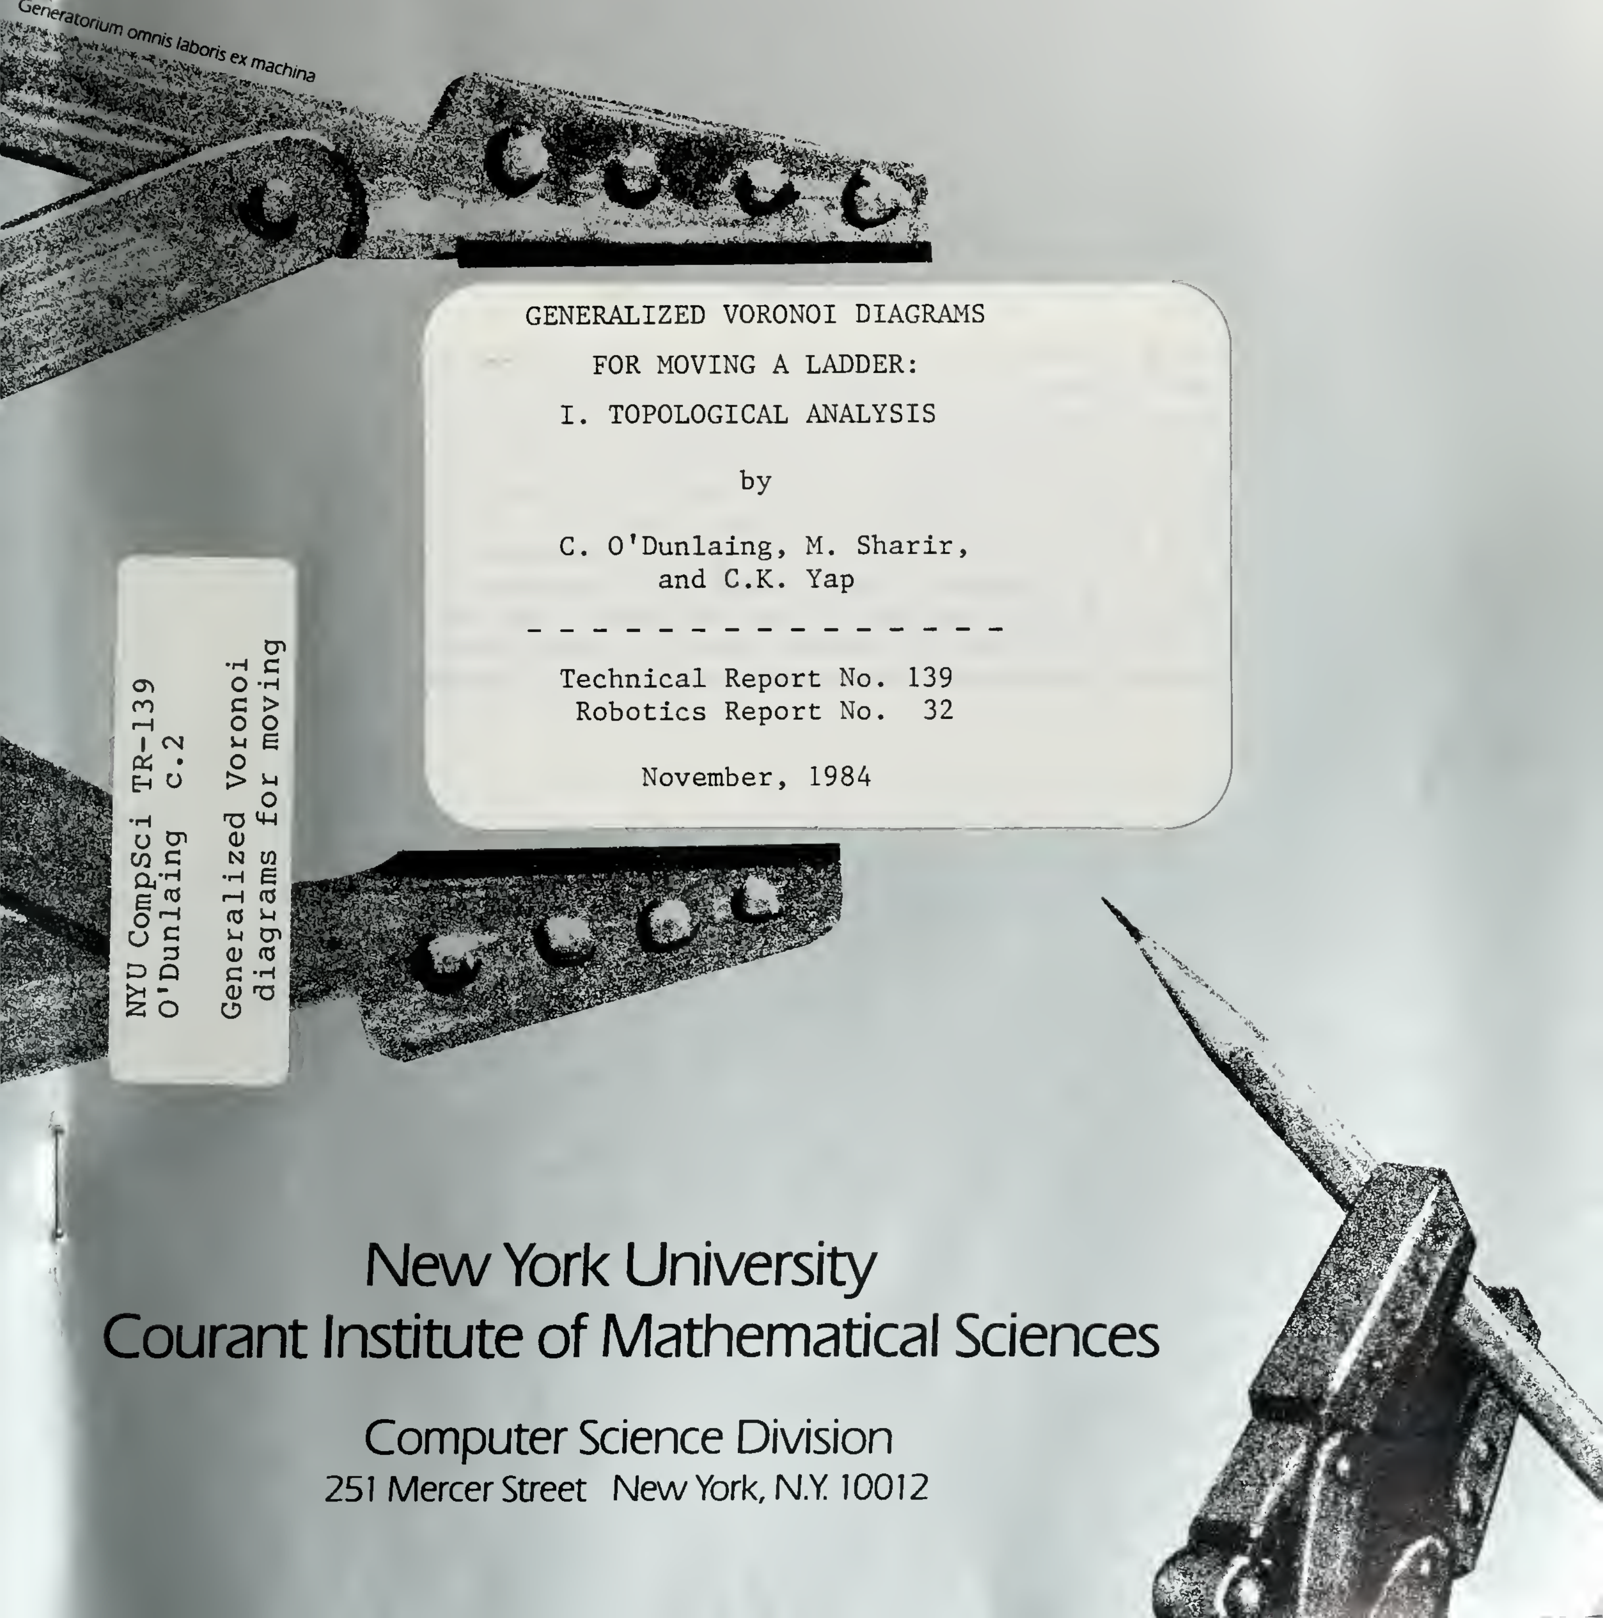

GENERALIZED VORONOI DIAGRAMS FOR MOVING A LADDER:

I. TOPOLOGICAL ANALYSIS

by

C. O'Dunlaing, M. Sharir,
and C.K. Yap

Technical Report No. 139
Robotics Report No. 32

November, 1984

NYU CompSci TR-139
O'Dunlaing c.2
Generalized Voronoi
diagrams for moving

New York University
Courant Institute of Mathematical Sciences

Computer Science Division
251 Mercer Street New York, N.Y. 10012

GENERALIZED VORONOI DIAGRAMS
FOR MOVING A LADDER:
I. TOPOLOGICAL ANALYSIS

by

C. O'Dunlaing, M. Sharir,
and C.K. Yap

Technical Report No. 139
Robotics Report No. 32

November, 1984

2

Generalized Voronoi diagrams for moving a ladder: I. topological analysis†

Colm Ó'Dúnlaing
*Micha Sharir**
Chee K. Yap

Courant Institute of Mathematical Sciences
New York University
251 Mercer Street
New York, New York 10012

ABSTRACT

Given a bounded open subset Ω of the plane whose boundary is the union of finitely many polygons, and a real number $d > 0$, a manifold *FP* (the 'free placements') may be defined as the set of placements of a closed oriented line-segment B (a 'ladder') of length d inside Ω. *FP* is a three-dimensional manifold. A 'Voronoi complex' in this manifold, a two-dimensional cell complex, is defined by analogy with the classical geometric construction in the plane; within this complex a 1-dimensional subcomplex N, called the *skeleton*, is defined. It is shown that every component of *FP* contains a unique component of N, and canonical motions are given to move the ladder to placements within N. In this way, general motion planning is reduced to searching in a suitable represention of N as a (combinatorial) graph. Efficient construction of N is described in a companion paper.

November 2, 1984

---

†This work has been supported in part by ONR grants N00014-75-C-0571 and N00014-82-K-0381, and by grants from the Digital Equipment Corporation, the Sloan Foundation, the System Development Foundation, and the IBM Corporation. Work by the second author has also been supported by the U.S.-Israeli Binational Science Foundation. Work by the other two authors has also been supported by NSF grant DCR-84-01898.

*Second author's permanent address: School of Mathematical Sciences, Tel-Aviv University, Ramat-Aviv, 69978 Tel-Aviv, Israel.

Generalized Voronoi diagrams for moving a ladder: I. topological analysis†

Colm Ó'Dúnlaing
*Micha Sharir**
Chee K. Yap

Courant Institute of Mathematical Sciences
New York University
251 Mercer Street
New York, New York 10012

Introduction.

Given a bounded region Ω (a 'room'; the boundary $\partial\Omega$ is assumed to he a finite union of polygons) in the plane, and a directed closed line-segment B (a 'ladder') of fixed length d but free to move within Ω, the *ladder-movers' problem* is to find, given a pair Z_0, Z_1 of placements of B within Ω, whether it is possible to move the ladder continuously between the two placements while remaining within Ω, and to plan such a motion if possible (since Ω is open, it is implicit that the ladder never touch the boundary of the room). Topologically, the set FP of 'free placements' of the ladder – i.e., placements within Ω – is an open three-dimensional manifold, and the movers' problem reduces to determining the (path-) connected components of this manifold. Our method is to identify a 'Voronoi complex' $Vor_d(\Omega)$ within the manifold FP, by analogy with the 'Voronoi diagram' in the plane, and further to define a 1-dimensional 'skeleton' N within this Voronoi complex. First we show that every (path-connected) component of FP contains a unique component of $Vor_d(\Omega)$. To show this we define a retraction-like mapping Im from FP onto $Vor_d(\Omega)$: if this map were continuous everywhere then the connectivity property would follow immediately; however, the map Im does possess discontinuities, but these are easily analyzed and it is shown that the discontinuities do not jump across different components of $Vor_d(\Omega)$.

The task of this paper is to investigate the topological properties of the Voronoi complex. To be precise, we separate the complex into connected two-dimensional surfaces with boundary, which we call *sheets*, where distinct sheets can overlap only along their boundaries; within each sheet we define a 1-dimensional connected network which includes the boundary of the sheet. The

---

†This work has been supported in part by ONR grants N00014-75-C-0571 and N00014-82-K-0381, and by grants from the Digital Equipment Corporation, the Sloan Foundation, the System Development Foundation, and the IBM Corporation. Work by the second author has also been supported by the U.S.-Israeli Binational Science Foundation. Work by the other two authors has also been supported by NSF grant DCR-84-01898.

*Second author's permanent address: School of Mathematical Sciences, Tel-Aviv University, Ramat-Aviv, 69978 Tel-Aviv, Israel.

union of all these networks defines the skeleton of FP mentioned above. The size of the skeleton as a combinatorial object is shown in a companion paper [OSY2] to be $O(n^2\log^* n)$, where n measures the complexity (number of corners) of the polygonally bounded room Ω. Construction of the skeleton in time $O(n^2\log n\log^* n)$ is also described in [OSY2].

There has been much activity recently on various forms of the movers' problem from a purely algorithmic point of view([LW],[SS1],[SS2],[HJW]). For the problem considered here of moving a ladder in Ω an $O(n^5)$ solution is described in [SS1], using a different approach. The use of the Voronoi diagram in motion-planning was first proposed in [Row] and subsequently rediscovered in [OY] where an efficient algorithm for moving a disc is described. An extended summary of the work in this paper (as well as that of [OY],[OSY2]) was first sketched out in the conference paper [OSY1].

1. Definitions and terms.

Let E_2 represent the Euclidean plane with the Euclidean metric. A 'ladder' B is a closed directed line-segment of fixed length d, free to move in a 'room' Ω – an open bounded region in E_2 whose boundary $\partial\Omega$ is a finite union of pairwise disjoint closed polygons. Following [Kir] we decompose $\partial\Omega$ into a finite set S of *corners* and open *wall-segments*. Here $n = |S|$ measures the complexity of the room. Any wall or corner in S is called an *object*. We regard each object as a set of points in the plane – thus a corner is a singleton set $\{p\}$. We assume for computational purposes that the set S is represented so that all wall-segments are given an explicit orientation, with the interior of Ω to their left. Thus there are no difficulties in regarding Ω as being completely defined by specifying S.

P and Q denote the two ends of the ladder. Each *placement* Z of the ladder can be specified as a triple (x,y,θ), where (x,y) are the co-ordinates of the P-end of the ladder and θ is the *orientation* of the ladder – the angle the directed line-segment PQ makes with the positive x-axis. Thus the set of placements has three degrees of freedom. Given a placement Z we write $P(Z)$, $Q(Z)$, and $\theta(Z)$ to indicate respectively the positions of the ends of the ladder and its orientation. Given a placement Z, the point set $B(Z)$ consists of all points on the ladder when in placement Z – to be precise, $B(x,y,\theta) = \{(x+t\cos\theta, y+t\sin\theta): 0 \le t \le d\}$. The placement Z is *free* if $B(Z)$ is in Ω. We are interested in the set $FP = FP(d,\Omega)$ of all free placements. Clearly it is an open submanifold of the set of all placements. The closure of FP we write as SFP and call the set of *semifree placements*. (Except for certain degenerate configurations of Ω, SFP is the set of all placements in which the ladder is in the closure of Ω.)

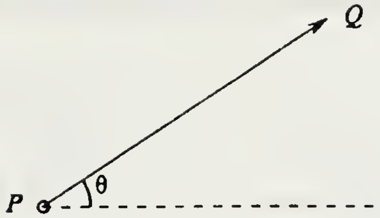

$Z = (x, y, \theta)$

$P(Z) = (x, y)$

$Q(Z) = (x + d\cos\theta,\ y + d\sin\theta)$

Figure 1. Illustrating a placement of the ladder

Given two closed nonempty sets X and Y in E_2, we define their separation, $sep(X,Y)$, as follows:

$$sep(X,Y) = Inf\ \{d(p,q): p \in X,\ q \in Y\}$$

and by abuse of notation write $sep(p,Y)$ for $sep(\{p\},Y)$, etcetera. Given a placement Z and a nonnegative real number r, the (Z,r)-*racetrack*, written $R(Z,r)$, is the planar set of points

$$\{p \in E_2: sep(p,Y) \le r\}$$

The shape of the 'racetrack' is the reason for its peculiar name. The boundary of the racetrack (assuming nondegenerate, i.e., $r > 0$), consists of two semicircles (the *bends*) and two straight line-segments (the *sides*). Given any placement Z, its *clearance*, written $Clearance(Z)$, is the smallest nonnegative r such that $R(Z,r)$ intersects the complement Ω^c of the room. (Alternatively, $Clearance(Z) = sep(B(Z),\Omega^c)$.) We write $R(Z)$, and call it the *Z-racetrack*, to denote the racetrack $R(Z,Clearance(Z))$. Note that FP can be characterized as the set of placements of positive clearance.

Given a semifree placement Z, we define $Near(Z)$ to be the intersection $R(Z) \cap \Omega^c$ (which is the same as $\partial R(Z) \cap \partial\Omega$). In placement Z an object s in S is *near* the ladder $B(Z)$ if $s \cap Near(Z) \neq \emptyset$. Finally, the *Voronoi complex* $Vor_d(\Omega)$ is the set of all free placements Z in FP such that $Near(Z)$ is *disconnected* (in general, $Near(Z)$ consists of a union of isolated points and closed line-segments). Note that this definition generalizes the *standard* notion of Voronoi diagram in the plane (see e.g. [OY]) in the sense that if $d = 0$ then the ladder degenerates to a point in the plane, FP may be identified with Ω, the racetracks become closed discs, and $Vor_0(\Omega)$ consists of all points Z in Ω such that $Near(Z)$ contains more than one point. Therefore we shall always write $Vor_0(\Omega)$ to denote this standard Voronoi diagram of Ω. A slightly different definition of $Vor_0(\Omega)$ is used by [Kir]. In particular, if s is a wall-segment and s' is its endpoint, then the Voronoi diagram of Kirkpatrick includes those points sufficiently close to s' on the line passing through s' perpendicular to s. We will call this alternative definition of Kirkpatrick the *augmented* standard Voronoi diagram. It is not hard to see that the augmented standard Voronoi diagram is precisely the union of the standard Voronoi diagram with additional edges of the type just described.

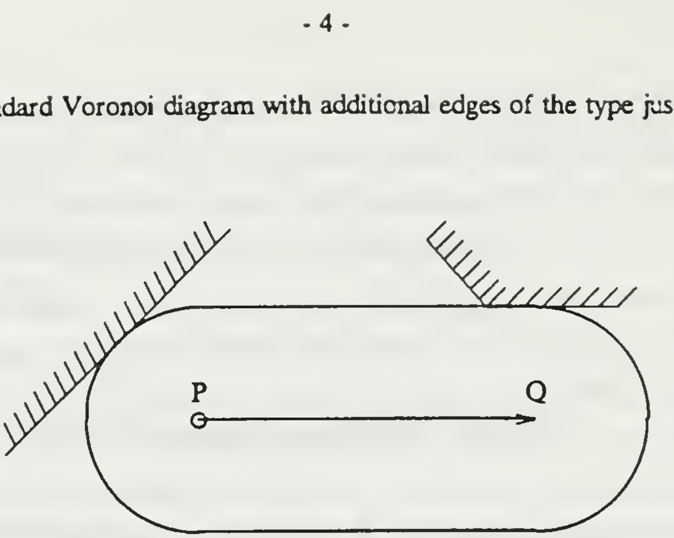

Figure 2. A Race-track: *Near(Z)* has two components

In analyzing FP and $Vor_d(\Omega)$ we shall make use of a 'displacement' technique which has proved useful in motion-planning. First, given two subsets X and Y of E_2, their *Minkowski sum* is

$$X+Y = \{p+q : p \in X,\ q \in Y\},$$

and similarly we define $X-Y$ as the set of differences $\{p-q : p \in X,\ q \in Y\}$. This should not be confused with the set-theoretic difference $X \sim Y$. Again by abuse of notation we write $p-X$ for $\{p\}-X$ and so on. Given an orientation θ, define B_θ as $B((0,0,\theta))$, i.e., the set of all points on the ladder when the P-endpoint is placed at (0,0) and the orientation is θ. Then define Ω_θ as

$$\Omega_\theta = (\Omega^c - B_\theta)^c$$

In other words, Ω_θ is the complement of the Minkowski sum of $-B_\theta$ with the complement of Ω. The next figure illustrates this.

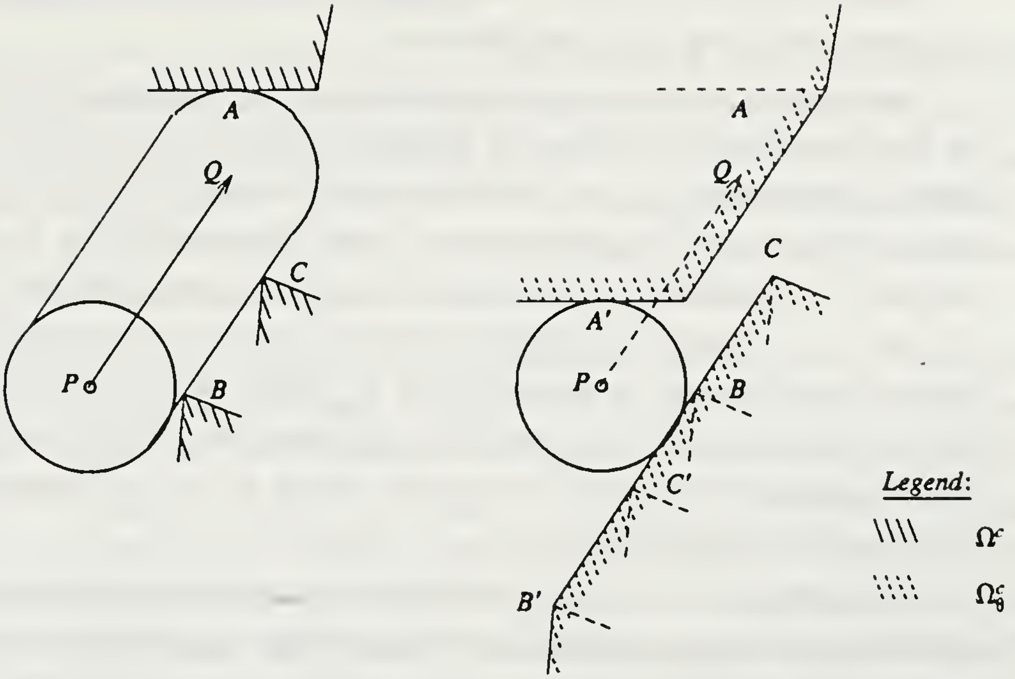

Figure 3. Illustrating Ω^c and Ω^c_θ

It is not difficult to verify the following relations:

[1] $B(Z) = P(Z) + B_{\theta(Z)}$.

[2] $R(Z,r) = D_r(x,y) + B_\theta$, where $Z = (x,y,\theta)$ and $D_r(x,y)$ is the closed disc of radius r about (x,y).

[3] $sep(p,B(Z)) = sep(P(Z),p-B_{\theta(Z)})$ for any point p in the plane. (Again note that $p-B_{\theta(Z)}$ is the Minkowski difference.)

[4] $Clearance(Z) = sep(P(Z),\Omega^c-B_\theta) = sep(P(Z),(\Omega_{\theta(Z)})^c)$.

A consequence of these relations is that we can loosely regard the Voronoi complex as a family of planar Voronoi diagrams parametrized by orientation θ: this is made precise in the next lemma. Recall that a free placement Z is in $Vor_d(\Omega)$ if $Near(Z)$ has more than one connected component, so at least two points in $\partial\Omega$ belong to different components of $Near(Z)$. We shall see later that some difficulties are introduced in considering placements in which $Near(Z)$ is disconnected but contained in the same side of the Z-racetrack; for instance, two corners of $\partial\Omega$ could touch the same side of the racetrack. We shall call placements in which this occurs *broadside* placements.

Lemma 1.1 Let $Z = (x,y,\theta)$ be a free placement. If $(x,y) \in Vor_0(\Omega_\theta)$, then $Z \in Vor_d(\Omega)$. The converse holds except for broadside placements.

Proof. Let $r = Clearance(Z)$, and suppose that (x,y) is in the standard Voronoi diagram of Ω_θ. Thus, by property [4] above and the definition of $Vor_0(\Omega_\theta)$, the closed disc D of radius r about (x,y) intersects $\partial\Omega_\theta$ in at least two points p and q. It is not difficult to deduce that $Near(Z)$ intersects both $p-B_\theta$ and $q-B_\theta$, at points p' and q', say. It is impossible that p' and q' be on the same side of the racetrack, so $Near(Z)$ has more than one component, and thus $Z \in Vor_d(\Omega)$. Conversely, if $Z \in Vor_d(\Omega)$ is not a broadside placement, then there exist at least two points p' and q' in $Near(Z)$ but not on the same side of the racetrack. It is easy to show that $D_r(Z)$ intersects both $p'-B_\theta$ and $q'-B_\theta$, where $D_r(Z)$ is the closed disc of radius $Clearance(Z)$ about $P(Z)$; since the sets $p'-B_\theta$ and $q'-B_\theta$ are clearly disjoint, the points of intersection must be distinct and $P(Z) \in Vor_0(\Omega_\theta)$. Q.E.D.

Let s_1 and s_2 be two objects in the set S of corners and wall-segments defining Ω. We say that the objects are *separated* unless either (*i*) they consist of a corner and an adjacent wall-segment, or (*ii*) they are two corners at opposite ends of a common wall-segment. It is easy to verify that a free placement Z is in $Vor_d(\Omega)$ if and only if $Near(Z)$ intersects two separated objects. Therefore, when two objects s_1, s_2 are separated we define the *midsurface* $M(s_1,s_2)$ as the set

$$\{Z \in FP: Near(Z) \subseteq s_1 \cup s_2,\ Near(Z) \cap s_1 \neq \emptyset,\ Near(Z) \cap s_2 \neq \emptyset\}.$$

In other words, $M(s_1,s_2)$ consists of all free placements near to s_1 and s_2 but to no other object in S. It is not hard to see that a midsurface need not be connected:

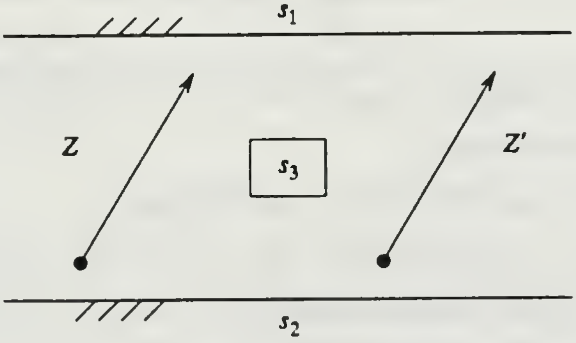

Figure 4. Z and Z' are in different components of $M(s_1, s_2)$

An (s_1,s_2)-*sheet* is defined as the closure in *SFP* of a connected component of $M(s_1,s_2)$. The *interior* K° of a sheet K is the set of all placements Z in K such that for some neighborhood U of Z (in the space of all placements, free or otherwise) $U \cap SFP \subseteq K$; the boundary ∂K is $K \sim K^\circ$; and the *free boundary* of K is the intersection $\partial K \cap FP$.

Remark. We have defined a sheet only for separated pairs of objects, since placements in $Vor_d(\Omega)$ must be near two or more separated objects. However, the definition of a sheet will force us to consider placements in which the ladder is near three or more objects which are not necessarily pairwise separated. Note also that if s_1 is a *concave corner*, i.e., one at which the angle of the incident walls as measured from inside Ω is less than π, then $M(s_1,s_2)$ is empty for all objects s_2. Most discussions in this paper become vacuous for concave corners.

We shall see in later sections that two sheets can intersect only along their boundaries, and that $Vor_d(\Omega)$ is covered by these sheets (the inclusion is strict since sheets can include semifree placements). Indeed, once the basic properties of $Vor_d(\Omega)$ have been established, the remaining analysis will investigate the properties of sheets. The correctness of the analysis will depend heavily on certain nondegeneracy conditions holding for FP. For example, in the figure below, the ladder is in $Vor_d(\Omega)$ but not on any sheet. This anomaly arises from the collinearity of the three corners shown in the figure (the placement is on no sheet since it is impossible for nearby placements to be near to two of the corners without being near to the third).

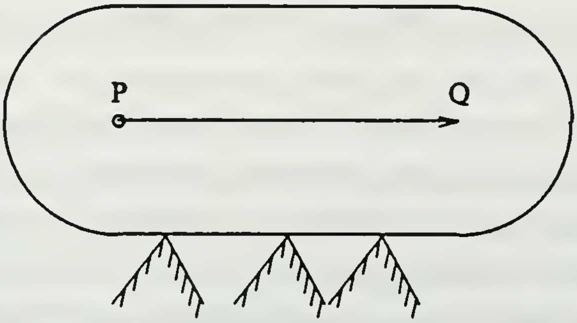

Figure 5. A placement not in any sheet

To eliminate this and other similar problems, we make the following nondegeneracy assumptions about Ω and the relation of the ladder to Ω.

> **General position.** No three corners are collinear; no four corners are concyclic; no two corners are separated by a distance d; no corner and a wall are separated by a distance d; no two pairs of corners determine parallel lines (whether or not the lines define wall-segments); for any free placement Z, $Near(Z)$ has at most 4 components; given any placement Z such that $B(Z) \subset \overline{\Omega}$ (closure of Ω), $Z \in SFP$ (in other words, there are no placements in which the ladder is 'stuck'); given two walls, the infinite lines containing them do not intersect at a point equidistant from any two corners.

Thick walls. Ω is the interior of its closure, with no slits or punctures; every corner is incident to exactly two walls.

Discussion. It is worthwhile at this stage to consider the broad strategy of the rest of the paper. In Section 2 we show that motion-planning in *FP* can be reduced to motion-planning in the Voronoi complex $Vor_d(\Omega)$. The idea is to define a map $Im: FP \to Vor_d(\Omega)$ which is retraction-like in the sense that the map fixes all placements in $Vor_d(\Omega)$ and is continuous almost everywhere. Essentially the map *Im* is defined by the rule that if a placement Z is not in $Vor_d(\Omega)$ then an image placement $Im(Z)$ in $Vor_d(\Omega)$ is obtained by translating the ladder (keeping its orientation fixed) in a direction of maximally increasing clearance until a placement in $Vor_d(\Omega)$ is obtained. It follows that Z and $Im(Z)$ must always be in the same (path)-component of *FP*. If Z_0 and Z_1 are placements in the same (path)-component of *FP* then by definition there is a continuous map $\Pi: [0,1] \to FP$ taking 0 to Z_0 and 1 to Z_1. The composition of the two maps *Im* and Π yields a map Φ from [0,1] to $Vor_d(\Omega)$. We shall see that although the map Φ may have discontinuities they can be assumed to be finite in number and then shown not to 'jump' across different path components of $Vor_d(\Omega)$. Thus a path from Z_0 to Z_1 can be used to define a path from $Im(Z_0)$ to $Im(Z_1)$ in $Vor_d(\Omega)$. This essentially reduces the problem of planning motions in *FP* to planning motions in $Vor_d(\Omega)$. (See [OY] for an application of the same principle to the simpler problem of moving a disc.)

In Section 3 we introduce standard choices of coordinate system for $M(s_1,s_2)$ for each of the different types of pairs of objects: two corners, two walls, or a wall and a corner. We also introduce the idea of 'streamlines' which are cross-sections of such midsurfaces obtained by fixing one of the parameters. (Specifically, except for the case of broadside placements, these streamlines are sets of placements all having a common orientation θ, so we have a family of streamlines each parametrized by orientation θ). In Section 4 we analyze the boundary of a sheet and classify it into several different classes depending on the relations which locally determine the boundary. In Section 5 we show that a streamline must be connected and must intersect the interior of the sheet in a connected interval of placements. It follows that each streamline intersects the boundary in two connected intervals of placements, which we call the (east and west) 'ends' of the streamline. The east (resp. west) end of a sheet is defined as the union of the east ends of streamlines in the sheet. This enables us to describe the boundary ∂K of a sheet K as the union of the east and west ends of K. In Section 6 we show that in any sheet K either (*a*) all orientations are feasible, and the sheet is homeomorphic to a bounded cylinder, with disconnected boundary, or (*b*) not all orientations are feasible within the sheet and the sheet is homeomorphic to a disc, with connected boundary. Based on this analysis we add one or two 'canonical' streamlines to the boundary of the sheet so to produce a connected 1-dimensional network containing the boundary of the sheet.

This we call the 'skeleton' of the sheet.

By taking the union of the skeletons of all sheets of $Vor_d(\Omega)$ we obtain a 1-dimensional network whose connectivity reflects the connectivity of $Vor_d(\Omega)$ and hence of FP. By this is meant that if there exists a path from Z_0 to Z_1 in $Vor_d(\Omega)$, then there exists a path which, except for its initial and final segments, is entirely within network. Computing the initial (resp. final) segment of such a path merely involves following the streamline containing the initial (resp. final) placement in a direction of nondecreasing clearance (this can always be done) until one or other end of the streamline is achieved. Since the connectivity of $Vor_d(\Omega)$ is fully defined by the adjacency relationships among its sheets, and therefore by the adjacency relation among the skeletons of its sheets, it follows that to find the rest of the path reduces to a combinatorial search of the skeleton of the sheet.

Finally, (Section 7), we investigate more closely the structure of the skeleton of a sheet and conclude that the *free skeleton* of every sheet is connected. This is convenient since it shows that the method of this paper yields motions which do not involve semifree placements. Without this observation the motions prescribed by simply following the (semifree) skeleton might need to be perturbed to contain them within FP.

Let us briefly compare the approach of this paper with the solution [OY] to the simpler problem of moving a disc. The latter problem involves computing the standard Voronoi diagram $Vor_0(\Omega)$, considering a map *Im* as before, and deducing that the motion-planning problem essentially reduces to one in which the center of the disc moves along the edges of the Voronoi diagram. The structure of the standard Voronoi diagram is simplified [Kir] if for each convex corner we add two edges emanating from the corner and perpendicular to the two walls incident to the corner. The augmented diagram is still a planar graph G with size linear in the number of objects; Also, the components of $\Omega \sim G$ (ie., the faces of the graph) are called Voronoi cells where each cell $V(s)$ is 'owned' by an object s in S in the sense that $V(s)$ consists of all points in Ω closer to s than to any other object in S. Each cell $V(s)$ is star-shaped relative to s and hence connected, and the map *Im* projects $V(s)$ onto its 'real' boundary – its boundary within the standard Voronoi diagram. The cells $V(s)$ are bounded by Voronoi edges which in turn are bounded by Voronoi vertices. In this paper the sheets K play the roles of such cells. Corresponding to Voronoi edges we define 1-dimensional sets in $Vor_d(\Omega)$ consisting of all those placements that are closest (simultaneously) to three object (s_1, s_2, s_3). We call these sets *midlines* and denote them as $M(s_1,s_2,s_3)$. Similarly the set of placements closest to four objects s_i define further subsets which by applying the principle of general position and using ideas of algebraic geometry we know to be finite sets: hence their elements are isolated placements which we also call *Voronoi vertices*.

Let us briefly consider two examples to show how the boundary of a sheet can be disconnected. In the first the ladder is trapped in a large triangular room bounded by three walls s_i. Here the ladder can be rotated completely while maintaining contact with s_1 and s_2 (or with one end touching their common corner). We deduce that one 'end' of the sheet $M(s_1, s_2)$ consists entirely of semifree placements, and is continuously parametrized by orientation θ, thus homeomorphic to the unit circle S^1. We know that every θ-streamline contains boundary placements at the other 'end' of the sheet (where the ladder is also nearest to s_3), and this set of boundary placements is (for the example illustrated here) also continuously parametrized by orientation θ. Thus in this example the boundary is disconnected but the free boundary (consisting entirely of the other ends of these streamlines) is connected.

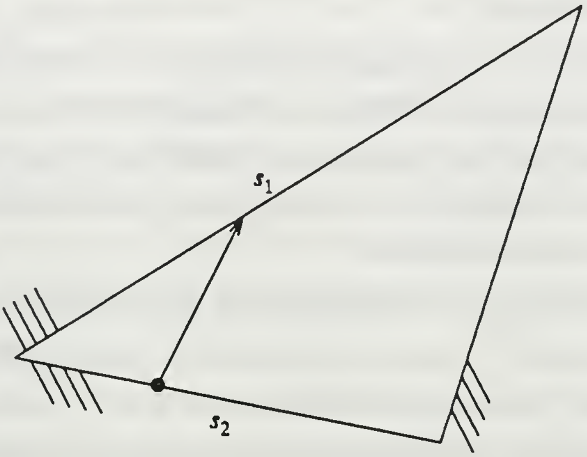

Figure 6. Sheet with disconnected boundary but its free boundary is connected

In the second example s_1 is a corner and s_2 is a horizontal wall whose distance from the corner is less than d, and s_1 has two incident walls s_3 and s_4 forming angles $\theta_{min} < 0$ and $\theta_{max} > 0$ respectively (see next figure). It can be shown that $[\theta_{min}, \theta_{max}]$ is the range of orientations feasible in one of the sheets of $M(s_1, s_2)$. In this case there exists one negative angle $\theta_0 > \theta_{min}$ at which P touches s_1 and Q touches s_2 and a positive angle $\theta_1 < \theta_{max}$ at which P touches s_2 and Q touches s_1. (See figure 20 in Section 7 for a similar configuration.) It is not hard to see that the boundary (composed of the east and west ends) of the sheet are connected. There is a connected subset of the boundary corresponding to semifree placements where P touches s_2 and s_1 touches Q or the interior of the ladder; similarly, another such subset corresponds to where Q touches s_2 and s_1 touches P or the interior of the ladder. The orientations of placements in these two semifree subsets lie in the range $[\theta_{min}, \theta_0]$ and $[\theta_1, \theta_{max}]$. Thus the boundary is connected but the free boundary has two components. [Remark: in forming the skeleton discussed above, in this case we augment the boundary with the set of all placements in K at orientation 0, and this is a connected set of free placements intersecting the east and west ends in free placements. This set serves to connect the components of the free boundary of K. Note that $M(s_1,s_2)$ has two components in this case; the other component corresponds to the range $[-\theta_{max}, -\theta_{min}]$.]

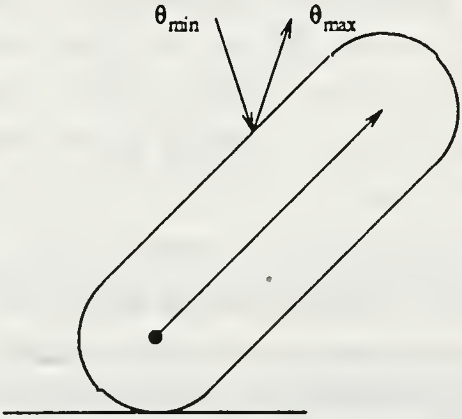

Figure 7. The boundary is connected but the free boundary is not

2. The retraction theorem.

In this section we show that every (path-connected) component of FP contains exactly one component of $Vor_d(\Omega)$, and furthermore that there is a retraction-like map Im which takes every placement Z to a placement $Im(Z)$ on $Vor_d(\Omega)$ and in the same component of FP. We begin with a basic property of the Voronoi complex of FP: it is closed in FP.

Definition. Let A and B be two nonempty subsets of a metric space Y. The *Hausdorff distance* $d(A,B)$ is defined as $\max(f(A,B),f(B,A))$ where

$$f(A,B) = Sup \{ sep(y,A) : y \in B \}$$

The Hausdorff distance defines a semi-metric on the set of all nonempty subsets of Y, and the restriction of the Hausdorff distance to the set of closed subsets of Y is a metric on that set. If $F: X \to Pow(Y) \sim \{\emptyset\}$ is a family of nonempty subsets indexed (or parametrized) by indices in a topological space X, we say that F is *continuous* if it is continuous with respect to the Hausdorff distance.

Remark. When in the above definition the space Y is Euclidean (in this paper Y is always the plane E_2 or a subset thereof), the map F is continuous if and only if for every z in X and $\epsilon >$ 0 there exists an open neighborhood U of z in X such that for every z' in U,

$$F(z) \subseteq F(z')+D_\epsilon ,$$

and

$$F(z') \subseteq F(z)+D_\epsilon ,$$

where D_ϵ is the (closed) disc of radius ϵ about (0,0). (Note that once again we are using the Minkowski sum of sets). This definition of continuity is easier to deal with. In order to show that $Vor_d(\Omega)$ is closed in FP we shall use the fact (proved below) that the Z-racetrack varies continuously with Z. Note that when we speak of SFP as a topological space, we are implicitly giving it, say, the relative topology as a subset of the metric space $E_2 \times S^1$ (S^1 is the unit circle). Loosely speaking, (x,y,θ) is close to (x',y',θ') if the components are close.

Lemma 2.1. The mapping $Z \rightarrow R(Z)$ is continuous on SFP.

Proof. (Note that the Z-racetrack coincides with $B(Z)$ if Z is not a free placement.) It is easy to show that the function $Clearance(Z)$ is continuous on SFP. Fix a placement Z in SFP, and choose $\epsilon > 0$. It is easy to show there exists a neighborhood U of Z in SFP such that for every Z' in U

(i) $|Clearance(Z) - Clearance(Z')| < \epsilon$,

(ii) $B(Z) \subseteq B(Z')+D_\epsilon$, and

(iii) $B(Z') \subseteq B(Z)+D_\epsilon$.

Therefore, writing r for $Clearance(Z)$, and r' for $Clearance(Z')$, $R(Z) = R(Z,r) = D_r+B(Z) \subseteq D_{r'+\epsilon}+(B(Z')+D_\epsilon) = D_{r'}+B(Z')+D_\epsilon+D_\epsilon = R(Z')+D_{2\epsilon}$. The reverse inclusion is proved similarly. Q.E.D.

In Lemma 2.2 below we exploit the continuity of $R(Z)$ as a function of Z in a way which will recur many times in this paper, and which is therefore worth separate consideration. First let us classify the components of $Near(Z)$ for any semifree placement Z. Let N be a component of $Near(Z)$, so it is either a point or a closed line-segment. Thus N may be either (i) a corner; (ii) a single point on a wall; (iii) a line-segment (of length d) contained in a wall; (iv) a line-segment intersecting a wall and one of its endpoints; or (v) a line-segment consisting of a wall and both of its endpoints.

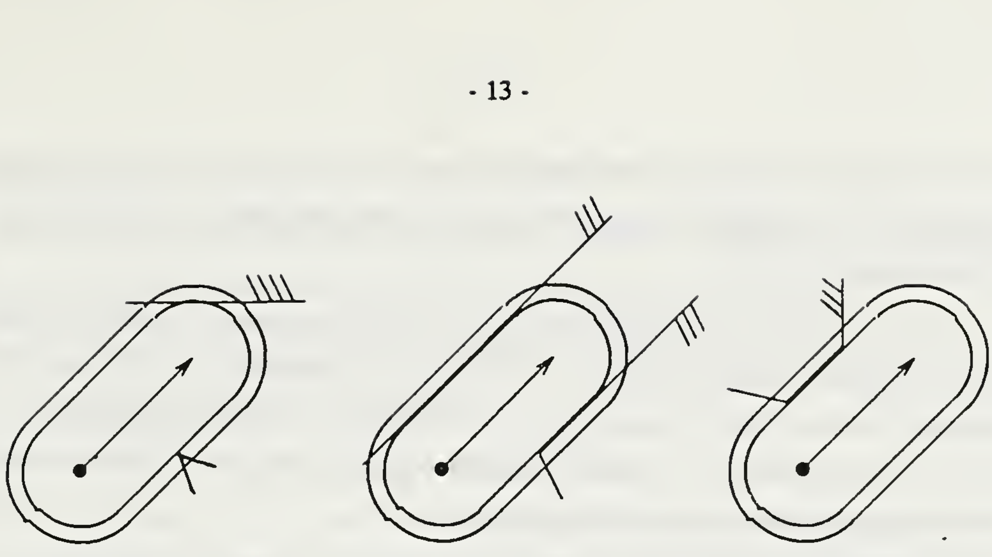
(i) and (ii)
(iii) and (iv)
(v)

Figure 8. Different types of components of *Near*(Z)

Definition. Let Z be a semifree placement. A *thickening* of the Z-racetrack is any racetrack R of the form $R(Z,r)$ where $r > Clearance(Z)$ but is sufficiently small to ensure that every connected component C of $R \cap \Omega^c$ is convex, hence connected, and contains exactly one component N of $Near(Z)$, and $C \cap \partial\Omega$ is contained in the following objects, depending on the above classification of N: (i) the corner and its two incident walls; (ii, iii) the wall itself; (iv) the corner and its two incident walls; (v) the two corners and the three incident walls. It is not hard to see that thickenings for any Z-racetrack exist. Note that this depends on the 'thick walls' assumption, which requires that every corner be incident to exactly two walls.

Corollary to Lemma 2.1. Let Z be a semifree placement. Then for any thickening R of the Z-racetrack there exists an open neighborhood U of Z in SFP such that for any placement Z' in $SFP \cap U$,

(i) $R(Z) \subseteq R$, and

(ii) every component of $R \cap \Omega^c$ contains at most one component of $Near(Z')$.

Proof. (i) is an easy consequence of the continuity of $R(Z)$ as a function of Z. To prove (ii), note that $Near(Z') = \bigcup (R(Z') \cap C_i)$ where the union is over all convex components C_i of $R \cap \Omega^c$. Q.E.D.

The above corollary will prove very useful; it readily implies that $Vor_d(\Omega)$ is closed in FP. (Since $Vor_d(\Omega) \subseteq FP$, clearly $Vor_d(\Omega)$ is not closed in SFP.)

Lemma 2.2. $Vor_d(\Omega)$ is closed in FP.

Proof. It is enough to show that $FP \sim Vor_d(\Omega)$ is open. Let Z be a placement in $FP \sim Vor_d(\Omega)$, and let R be a thickening of the Z-racetrack; by definition, $R \cap \Omega^c$ is convex. By the above corollary there exists an open neighborhood U of Z in the (open) set FP such that for every placement Z' in U, $Near(Z')$ has at most one component, and hence $Z' \in FP \sim Vor_d(\Omega)$. Thus $U \subseteq FP \sim Vor_d(\Omega)$. Q.E.D.

We now define a flow-like function $\phi(Z,t)$ which is guaranteed to move the ladder onto the Voronoi complex. In general this flow is defined for all placements in $FP \sim Vor_d(\Omega)$, but it is important to observe that it makes sense also to define it for broadside placements as well. This observation is important in reconciling the discontinuities of the map *Im* (defined below) with the continuity requirement in motion-planning paths.

Definitions and remarks. Let Z be a placement in $FP \sim Vor_d(\Omega)$, so $Near(Z)$ is connected. Note that $Near(Z)$ may be infinite, so there may be infinitely many points z in $B(Z)$ and q in $Near(Z)$ such that $d(q,z) = Clearance(Z)$; however, the difference $v_Z = z - q$ is unique and nonzero, and indeed coincides with $P(Z) - p$, where p is the unique point in $\partial\Omega_{\theta(Z)}$ closest to $P(Z)$. (The vector v_Z actually indicates the direction in which the ladder should be translated to increase clearance as rapidly as possible). For any t in **R** (the set of real numbers), define

$$P_t = P(Z) + tv_Z,$$

and

$$\phi(Z,t) = (P_t, \theta(Z)).$$

Note that if s_1 and s_2 are corners and Z is a broadside placement in $M(s_1,s_2)$ (so $Near(Z)$ is contained in a side of $R(Z)$), then the vector v_Z is also uniquely defined, and the 'flow' $\phi(Z,t)$ may again be defined unambiguously as above.

The following somewhat technical lemma is based on a simple intuition. Suppose that a placement Z is in $FP \sim Vor_d(\Omega)$. Then if we translate the ladder away from the closest obstacle as dictated by the 'flow' $\phi(Z,t)$, the ladder will recede from the closest obstacle until a time t_1 is reached at which the ladder attains the same clearance with another obstacle. The placement $Z_1 = \phi(Z,t_1)$ must be in $Vor_d(\Omega)$, and Z_1 is a placement at which the direction of maximally increasing clearance is either undefined or at least differs from v_Z. To show this we shall invoke the following simple observation: let L be the tangent to the racetrack $R(Z)$ containing $Near(Z)$, and for $t > -1$ let R_t be the racetrack around $\phi(Z,t)$ tangent to L. Clearly, $R_t \supseteq R(\phi(Z,t))$. Then the sets R_t are monotonically increasing with respect to set-inclusion, and indeed for any 'times' u and v such that $-1 < u < v$, any point in $R_u \sim L$ is interior to $R_v \sim L$, whereas $R_u \cap L = R_v \cap L$.

Lemma 2.3.

(i) For every placement Z_0 in $FP \sim Vor_d(\Omega)$ there is a unique real number $t(Z_0)>0$ such that

(*a*) For every t in the interval $(-1, t(Z_0))$, $Near(\phi(Z_0,t)) = Near(Z_0)$;

(*b*) $\phi(Z_0,t(Z_0)) \in Vor_d(\Omega)$, and this placement has locally maximal clearance among all placements of the form $\phi(Z_0,t)$; and

(*c*) $Near(Z_0)\cap Near(\phi(Z_0,t)=\emptyset$ for every $t>t(Z_0)$.

(ii) The same results hold if Z_0 is a broadside placement in $M(s_1,s_2)$ when s_i are both corners, with the following stronger formulation of (ii.*b*): $Near(\phi(Z_0,t(Z_0)))$ intersects some third object s_3 which is separated from s_1 and s_2.

Proof. The same proof holds for both (i) and (ii), so we shall treat them simultaneously. Let L be the line containing $Near(Z_0)$ and tangent to $R(Z_0)$ as introduced in the above discussion, and define the racetracks R_t as above. The racetracks R_t have radius $(1+t)Clearance(Z_0)$ and they expand in the manner described above. Hence (since Ω is bounded) there exists a 'time' $t(Z_0)$ such that $R_{t(Z_0)}$ touches $\partial\Omega$ at some point q not in L, but contains no point of Ω^c in its interior. Clearly, $R_t \cap L = R(Z_0)\cap L$ for all $t > -1$; thus given any t such that $-1<t<t(Z_0)$, $R_t\cap\Omega^c \subseteq R_t \cap L = R(Z_0)\cap L$, and hence $Near(\phi(Z_0,t)) = Near(Z_0)$, proving (*a*). Clearly, by continuity, $Near(\phi(Z_0,t(Z_0))) \supseteq Near(Z_0)$, but q must be in a different connected component of $Near(\phi(Z_0,t(Z_0)))$, and hence must lie in some object s_3 not intersecting $Near(Z_0)$; in particular, the placement is in $Vor_d(\Omega)$. This proves (*b*). Since q is interior to R_t for all $t > t(Z_0)$, (*c*) is immediate: the ladder must be closer to q than to $Near(Z_0)$ for all such values of t. Q.E.D.

Let us now define the map Im. The above simple lemma has covered the only relevant technical detail in analyzing the discontinuities of this map.

Definition. Given a placement Z in FP, if $Z \in Vor_d(\Omega)$ then

$$Im(Z) = Z,$$

otherwise

$$Im(Z) = \phi(Z,t(Z)),$$

where $t(Z)$ is the minimum $t>0$ (identical to the corresponding number in Lemma 2.3) such that $\phi(Z,t) \in Vor_d(\Omega)$.

Lemma 2.4. The map Im is continuous on $FP \sim Vor_d(\Omega)$. The only discontinuities possible for this map are where a sequence Z_n converges to Z_0 in FP, Z_0 is a broadside placement, and the sequence $Im(Z_n)$ converges to a unique placement Y_0 of the form $\phi(Z_0,u)$ for some $u>0$ and the entire set $\{\phi(Z,t):0\le t\le u\}$ is in $Vor_d(\Omega)$.

Proof. Let Y_n denote $Im(Z_n)$. By passing to a subsequence if necessary we may assume that Y_n converge to a placement Y_0. If infinitely many placements Z_n are in $Vor_d(\Omega)$ then $Y_n = Z_n$ infinitely often, so $Y_0 = Z_0$. Since $Vor_d(\Omega)$ is closed in FP, Z_0 is in $Vor_d(\Omega)$ so $Z_0 = Im(Z_0) = Y_0$. Thus we may assume that all but finitely many placements Z_n, and hence, without loss of generality, all placements Z_n, are in $FP \sim Vor_d(\Omega)$. By passing to a further subsequence if necessary we can choose, for each n, points p_n, z_n, and y_n in $Near(Z_n)$, $B(Z_n)$, and $B(Y_n)$ respectively, converging to points p_0, z_0, and y_0, such that p_n, z_n, and y_n are collinear, $d(p_n,z_n) = Clearance(Z_n)$, and $d(p_n,y_n) = Clearance(Y_n)$. By continuity we know that these relations hold in the limit as well (i.e., for $n = 0$).

Again by continuity, $Clearance(Z_0) \le Clearance(Y_0)$ so Y_0 is in FP and therefore (since $Vor_d(\Omega)$ is closed in FP) in $Vor_d(\Omega)$, and $\theta(Z_0) = \theta(Y_0)$. It is easy to verify from the collinearity of p_0, z_0, and y_0, that for some $u \ge 0$ $Y_0 = (P(Z_0)+u(z_0-p_0),\theta(Z_0))$. If $u = 0$ then $Y_0 = Z_0 = Im(Z_0)$ as before. Let us therefore assume that $u>0$.

First suppose that $Z_0 \in FP \sim Vor_d(\Omega)$. Then clearly $\phi(Z_0,u)$ is defined and equals Y_0. Since $p_0 \in Near(Y_0) \cap Near(Z_0)$, it follows from Lemma 2.3 (i*c*) that $u \le t(Z_0)$, whereas, since $t(Z_0)$ is the smallest t such that $\phi(Z_0,t) \in Vor_d(\Omega)$, $u \ge t(Z_0)$; hence $u = t(Z_0)$ and $Y_0 = Im(Z_0)$ again. This shows that Im is continuous on $FP \sim Vor_d(\Omega)$.

Thus we need only consider the case where $Z_0 \in Vor_d(\Omega)$. By applying the methods of Lemma 2.3 it is easy to show that $Near(Z_0)$ must be contained in a straight line L, since otherwise (since $u>0$) the racetrack around Y_0 and containing p_0 would contain other boundary points in its interior. Hence Z_0 must be a broadside placement; thus, again, $Y_0 = \phi(Z_0,u)$. Since $p_0 \in Near(Z_0) \cap Near(Y_0)$, it follows from Lemma 2.3 (ii*c*) that $u \le t(Z_0)$. It follows (by Lemma 2.3 (ii*a*)) that $\{\phi(Z_0,t)\colon 0<t<u\}$ is a path connecting Z_0 to Y_0 and lying entirely within $Vor_d(\Omega)$.

This proves, at any rate, that if Im is discontinuous at Z_0 then the latter placement is a broadside placement and for any limit point Y_0 of $Im(Z_n)$, Z_0 and Y_0 are in the same path-connected component of $Vor_d(\Omega)$. It is worth observing, however, that there is only one such limiting placement Y_0: to be precise, $Y_0 = \phi(Z_0,t(Z_0))$. (Recall that Y_0 was originally chosen by passing to a subsequence.) To show this it is enough to assume that Z_0 is a point of discontiuity and show that $u \ge t(Z_0)$, where u is the number defining Y_0 in terms of ϕ. Note in this case that $Near(Z_0) = s_1 \cup s_2$ consists of two corners (any other case contradicts the assumption of general position). Let R be a thickening of $R(Z_0)$. For all sufficiently large n, $R(Z_n) \subseteq R$, and $Near(Z_n) = R \cap Near(Y_n)$; therefore (by passing to a subsequence if necessary) there exists a sequence of points q_n in $Near(Y_n) \sim R$ converging to some point q_0. By continuity, $q_0 \in Near(Y_0) \sim R$. Thus since $Y_0 = \phi(Z_0,u)$, we know (by Lemma 2.3, (ii*a*)) that $u_0 \ge t(Z_0)$, proving the assertion. Q.E.D.

Remark. An obvious but crucial corollary to the above lemma is that if *Im* is discontinuous at Z_0 then there are only two limiting placements for $Im(Z)$ near Z_0, namely Y_0 and Z_0 itself, and both these placements are in the same path-connected component of $Vor_d(\Omega)$.

Theorem 2.5 (the 'retraction theorem.') For any free placement Z, Z and $Im(Z)$ are in the same (path-connected) component of FP. Furthermore, two placements Z_0 and Z_1 are in the same component of FP if and only if $Im(Z_0)$ and $Im(Z_1)$ are in the same component of $Vor_d(\Omega)$.

Proof. If $Z \neq Im(Z)$ then there is a clearance-increasing path from Z to $Im(Z)$ parametrized in the form $\phi(Z,t)$: thus Z and $Im(Z)$ are in the same path-connected component of FP. Therefore, if $Im(Z_i)$ are in the same path-connected component ($i = 0,1$), then so are Z_i.

Given a path, say $\{Z_t : 0 \leq t \leq 1\}$ from Z_0 to Z_1 in FP, we may by an argument based on Sard's Lemma (as in [SS1]) assume that the path intersects $Vor_d(\Omega)$ only finitely often. The path $Im(Z_t)$ is continuous except perhaps where the path intersects $Vor_d(\Omega)$. However, by the remark following Lemma 2.4, if s represents a point of discontinuity so $Im(Z_s) \neq Im(Z_{s+})$, say, then these two placements are still in the same path-connected component of $Vor_d(\Omega)$. Therefore $Im(Z_0)$ and $Im(Z_1)$ are in the same path-connected component of $Vor_d(\Omega)$. Q.E.D.

The above proposition shows us how to reduce the number of degrees of freedom of the movers' problem to two. Notice that given a placement Z, construction of $Im(Z)$ is a simple task computationally - by considering each corner and wall-segment in turn one can compute $Near(Z)$ in time $O(n)$, and it then becomes a simple matter to compute $Im(Z)$ in time $O(n)$. (Recall that $n = |S|$.)

3. Standard parametrizations of a sheet.

Given two separated objects s_1 and s_2, recall that an (s_1,s_2)-sheet is the closure of a (path)-connected component of $M(s_1,s_2)$, where the midsurface $M(s_1,s_2)$ itself consists of all placements in $Vor_d(\Omega)$ near to s_1 and s_2 but to no other object. In this section we discuss convenient ways of parametrizing $M(s_1,s_2)$ so we can further study the properties of sheets.

At this point, it is helpful to consider what further analysis needs to be done. It is necessary (i) to show that every placement in $Vor_d(\Omega)$ lies on some sheet; (ii) to characterize the boundary of a sheet, analyzing both the semifree part and the free part; and (iii) to construct a 1-dimensional network in every sheet which is connected and which contains the free boundary of a sheet. Notice that in order to investigate the boundary of a sheet it is necessary to characterize its interior (relative to $Vor_d(\Omega)$). It is not even certain at this stage that the interiors of two sheets could not intersect. (We shall see that this cannot happen.)

Definition. Let s_1 and s_2 be two separated objects. First we change the coordinate system, if necessary, so that they are in the following *standard positions*, depending on their type:

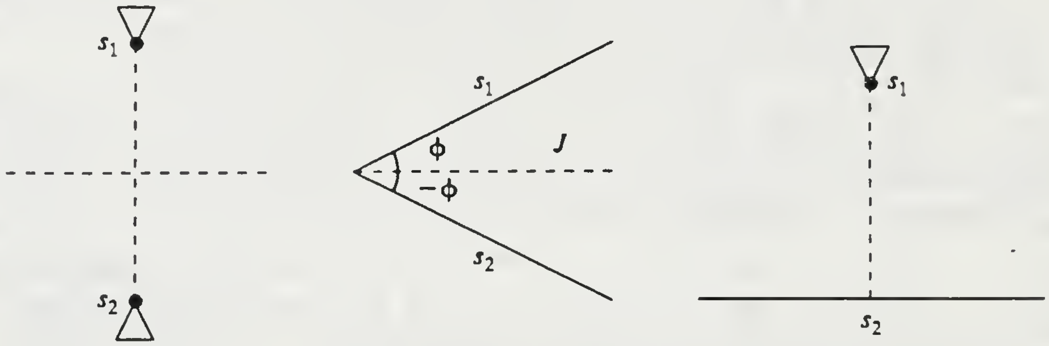

[CC] Case

[WW] Case

[WC] Case

Figure 9. The standard positions

[CC] If s_1 and s_2 are both corners, we assume that the axes are chosen so they have coordinates $(0,h)$ and $(0,-h)$ respectively. In other words, we assume they are on the y-axis and the x-axis is the locus of points equidistant from them.

[WW] If s_1 and s_2 are both walls, for $i = 1,2$ let H_i be the closed half-plane whose boundary contains s_i and such that near s_i, H_i and Ω are on the same side of s_i. These two half-planes must intersect (by general position) in a 'wedge' J. We orient the coordinate system so that the apex of the wedge J is at the origin, and the x-axis bisects the wedge, forming angles $\pm\phi$ $(0 < \phi < \pi/2)$ with the two rays bounding the wedge, and so that s_1 (respectively, s_2) intersects the ray extending at angle ϕ (respectively, $-\phi$) from the origin. Here the positive x-axis is the locus of points equidistant from the rays bounding the wedge. Note that for the ladder to be equidistant from s_1 and s_2 it is necessary that it be contained in the wedge J.

[WC] In the third case we assume that s_2 is a wall lying along the x-axis and $s_1 = \{(0,h)\}$ is a corner placed on the positive y-axis. In this case the locus of points equidistant from s_1 and the x-axis is a parabola with s_1 as focus and the x-axis as directrix. Note that for the ladder to be equidistant from s_1 and s_2 it is necessary that it be contained in the half-plane $\{y \geq 0\}$.

We call the rays bounding the wedge J in the [WW] case and the x-axis in the [WC] case *extended walls*. Now we define the *virtual midsurface* $VM(s_1,s_2)$ as the set of placements in which the ladder is equidistant from the corners [CC], or the two rays enclosing the walls [WW], or the corner and the x-axis [WC]. In this definition all objects apart from s_1 and s_2 are ignored, and we allow placements in which the ladder touches the corners or extended walls. Clearly in each case $VM(s_1,s_2)$ strictly includes $M(s_1,s_2)$, and indeed contains placements outside SFP. Also in the [WW] case, $VM(s_1,s_2)$ includes those placements of the ladder in the wedge where one end touches the origin. Again, we single out as a special case the broadside placements in the [CC] case - those placements in which the ladder is vertical.

Lemma 3.1 Except for broadside placements in the corner-corner case, $VM(s_1,s_2)$ is (continuously) parametrized by (x,θ); the sets of broadside placements for each of the two angles $\pm 1/2\pi$ are parametrized by (x,y).

Proof. A simple case analysis ●

The above-mentioned parametrizations are inconvenient when s_i are both corners, since in this case the domain of definition of parametrization for the broadside placements is not open in $VM(s_1,s_2)$. As it happens, however, both parametrizations usually 'work' in the sense that the mappings $(x,y,\theta)\rightarrow(x,y)$ and $(x,y,\theta)\rightarrow(x,\theta)$ are local homeomorphisms. The parametrizations generally fail where they are not surjective or not injective as functions (for example, the mapping $(x,y,\theta)\rightarrow(x,\theta)$ is not surjective near broadside placements). The following lemma explains roughly why this is so.

Lemma 3.2. Let Z_0 be a placement in $Vor_d(\Omega)$. If there exists an open neighborhood U of Z_0 in FP of the form $(x_1,x_2)\times(y_1,y_2)\times(\theta_1,\theta_2)$, written $V\times\Theta$, such that for every (x,y) in V there exists a unique θ in Θ such that $(x,y,\theta) \in Vor_d(\Omega)$, then (x,y) parametrizes $U\cap Vor_d(\Omega)$. If for every (x,θ) in $(x_1,x_2)\times\Theta$ there exists a unique y in (y_1,y_2) such that $(x,y,\theta) \in Vor_d(\Omega)$ then (x,θ) parametrizes $U\cap Vor_d(\Omega)$.

Proof. To prove the first part, it is enough to show that the mapping $(x,y)\rightarrow(x,y,\theta)$ is continuous, and this follows easily from the fact that $Vor_d(\Omega)$ is closed in FP. The second observation is essentially the same. ●

Corollary to Lemma 3.1. Suppose that s_1 and s_2 are both corners, and $Z_0 \in VM(s_1,s_2)$ is a free (virtual) broadside placement which is extremal in the sense that one of the corners is closest to an endpoint of the ladder. Then the (x,y) parametrization of the set of broadside placements in $VM(s_1,s_2)$ can be extended to include an open neighorhood of Z_0.

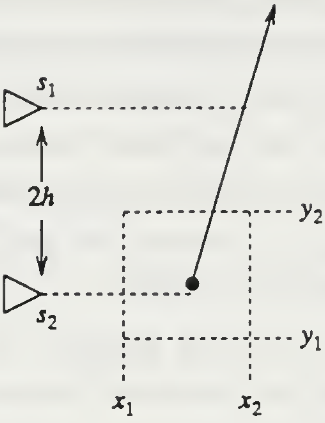

Figure 10. Parametrization about Z_0

Proof. Without loss of generality (choosing the standard coordinate-system), $Z_0 = (x_0, -h, \pi/2)$ where $x_0 > 0$ (in other words, we assume that the ladder lies east of the y-axis with its P-end closest to the lower corner s_2). Note (by general position) that s_1 must be closest to the interior of the ladder. Write $f(Z)$ for $sep(s_1, B(Z)) - sep(s_2, B(Z))$. It is easy to show that for a sufficiently small neighborhood U of Z_0 of the form $(x_1,x_2) \times (y_1,y_2) \times (\theta_1,\theta_2)$, for every point (x,y) in $(x_1,x_2) \times (y_1,y_2)$, $f(x,y,\theta_1) > 0$, $f(x,y,\theta_2) < 0$, and $f(x,y,\theta)$ has at most one (and hence exactly one) zero for θ in (θ_1,θ_2). Since f is clearly continuous and its zeroes define $VM(s_1,s_2)$, the corollary holds. Q.E.D.

We have now established convenient parametrizations – i.e., co-ordinate systems – for $VM(s_1,s_2)$ in all possible cases. Given an orientation θ (excluding orientations corresponding to broadside placements), the *virtual θ-streamline* consists of all placements (x,y,θ) within $VM(s_1,s_2)$. Thus each virtual θ-streamline is continuously parametrized by the x-coordinate. Similarly one speaks of virtual y-streamlines in the context of broadside placements: again these are continuously parametrized by x-coordinate. A virtual streamline can therefore be regarded as a parametrized family of placements Z_x where x range over the real numbers. The concept of virtual streamline will help us in analyzing the topology of an (s_1,s_2)-sheet K.

Definition. An *(s_1,s_2)-streamline* is the intersection of a virtual streamline with some (s_1,s_2)-sheet K.

The most significant fact about streamlines is that they are connected. (Connectivity of virtual streamlines is obvious.) This crucial property will be proved in a later section. The argument will generally be along the following lines: if we move along a virtual streamline past some point in ∂K, the ladder will move strictly closer to some other object in S. This is the spirit of the argument, but the precise details are much more complex (except for broadside placements in the [CC] case, where the proof is not very different from arguments in Section 2).

We begin by associating with placements Z of the ladder in $VM(s_1,s_2)$ certain sets of points which grow in a certain way as Z moves along a virtual streamline.

Definition. Let X_t be a family of subsets of the plane E_2 parametrized by a real parameter t. We say that the family is *expanding* if given $u < v$ for which X_u and X_v are defined,

$$X_u \subseteq X_v,$$

and say the expansion is *strict* except on some set Y if

$$X_u \sim Y \subseteq X_v^{\circ} \sim Y,$$

where W° denotes the interior of a set W. Similarly, we speak of the family as *contracting* (respectively, *strictly contracting*) with the obvious meaning. In the next set of definitions, when we speak of 'clearance' and 'racetrack', we mean 'virtual clearance' and 'virtual racetrack' in the sense that if s_1 (resp. s_2) is a wall, then the (virtual) clearance is defined as the distance of $B(Z)$ from the extended wall (the infinite line containing it), and all other objects are ignored.

Definitions. Let Z be a placement in $VM(s_1,s_2)$. We associate two or three planar sets with Z depending on the case being considered.

[CC] Define $\Gamma(Z)$ (respectively, $\Sigma(Z)$) to be the intersection of the (virtual) Z-racetrack with the half-plane $x \geq 0$ (respectively, $x \leq 0$). Thus $R(Z) = \Sigma(Z) \cup \Gamma(Z)$.

[WW] Recall that the ladder must be in a closed wedge J bounded by two rays emanating from the origin. If in placement Z one end of the ladder touches the apex, define $\Gamma(Z) = B(Z)$ and $\Sigma(Z) = J$. Otherwise, $J \sim R(Z)$ has two connected components $E(Z)$ and $W(Z)$ consisting of points east (respectively, west) of the racetrack $R(Z)$. Define $\Gamma(Z) = W(Z) \cup R(Z)$, $\Sigma(Z) = E(Z) \cup R(Z)$. Note that $\Gamma(Z)$ is bounded and $\Sigma(Z)$ is unbounded, and the virtual Z-racetrack is $R(Z) = \Sigma(Z) \cap \Gamma(Z)$.

[WC] Let A be the infinite strip $0 \leq y \leq h$ where $s_1 = \{(0,h)\}$. For any placement Z in $VM(s_1,s_2)$ $A \sim R(Z)$ has two connected components $E(Z)$ and $W(Z)$. Define $\Gamma(Z)$ and $\Sigma(Z)$ as $W(Z) \cup (A \cap R(Z))$ (respectively, $E(Z) \cup (A \cap R(Z))$), and define $\Delta(Z)$ as the intersection of $R(Z)$ with the half-plane $y \geq h$. Thus $R(Z) = (\Sigma(Z) \cap \Gamma(Z)) \cup \Delta(Z)$.

The significant properties of these sets is that $\Gamma(Z_\tau)$ is an expanding family and $\Sigma(Z_\tau)$ a contracting family as Z_τ moves eastward along a streamline. The behavior of $\Delta(Z)$ is unimodal: to see this, suppose s_1 is a corner and s_2 a wall; given a fixed orientation θ, if θ is not horizontal (i.e., $\theta \neq 0$ or π), then there is a unique placement Z_{τ_0} in the virtual θ-streamline at which the (virtual) clearance is minimal (perhaps zero). Say a placement Z_τ is *west* (respectively, *east*) of s_1

if $x \leq x_0$ (respectively, $x \geq x_0$). The following lemma summarizes the growth properties of the sets $\Gamma(Z)$, $\Sigma(Z)$, and $\Delta(Z)$. The proof of the lemma depends on a lengthy and tedious but simple geometric analysis and will be deferred to an appendix.

Lemma 3.3 (the 'monotonicity lemma.') The growth properties of the sets defined above, as $Z = Z_x$ moves along a streamline parametrized by the x-coordinate (so x is increasing), are, case by case, as follows:

[CC] The sets $\Gamma(Z)$ are expanding and the sets $\Sigma(Z)$ are contracting. The expansion is strict except where the racetrack intersects the y-axis, and on a side of the racetrack if that side touches one of the corners.

[WW] The sets $\Gamma(Z)$ are expanding and the sets $\Sigma(Z)$ are contracting. The expansion (contraction) is strict except on the boundary of the wedge J.

[WC] The sets $\Gamma(Z)$ are expanding and the sets $\Sigma(Z)$ are contracting. The expansion (contraction) is strict except on the boundary of the strip A and on a side of the racetrack if that side touches s_1. The sets $\Delta(Z_x)$ are contracting when Z_x is west of s_1 and expanding with Z_x is east of s_1; in the case where θ is horizontal and s_1 is closest to the interior of the ladder (so when $\theta = 0, -d < x < 0$, and when $\theta = \pi, 0 < x < d$), $R(Z_x) = \Gamma(Z_x) \cap \Sigma(Z_x)$. The contraction (expansion) is strict except on the line $y=h$ and on a side of the racetrack if that side touches s_1. ●

The above lemma will prove useful in relating facts about the local behavior of the racetrack $R(Z)$ when Z is in an (s_1,s_2)-sheet K to the 'global' behavior in $VM(s_1,s_2)$. We have already seen such observations about the 'virtual' racetrack in the discussions of Section 2. A more intuitive version of the monotonicity lemma may be expressed in a 'local' form. First observe that under the standard parametrizations when Z is a 'virtually' free placement, the boundary $\partial R(Z)$ of the (virtual) racetrack is separated by the two components of $Near(Z)$ into well-defined east and west segments.

Lemma 3.4 (local monotonicity lemma). Let p be a point on $\partial R(Z_u) \sim Near(Z_u)$ where Z_t is the standard parametrization (relative to s_1 and s_2) of the virtual (s_1,s_2)-streamline containing Z_u, where the virtual clearance of Z_u is nonzero. If p touches the eastern (respectively, western) segment of the boundary of $R(Z)$ then for all v sufficiently close to u and greater than (respectively, less than) u, p is in the (virtual) racetrack $R(Z_v)$. Furthermore, p is interior to $R(Z_v)$ unless it touches a side of the racetrack containing a component of $Near(Z_u)$.

Proof. A detailed analysis of the statement and proof of Lemma 3.3 will show that the local monotonicity properties follow from that lemma. ●

Finally, we state a basic result regarding the Σ, Γ and Δ sets.

Lemma 3.5 The families Σ, Γ, and Δ are continuously parametrized by Z.

The proof is similar to that for Lemma 2.1 and will be omitted. Note that continuity holds even near placements where the virtual clearance is zero.

4. Sheets: their interiors and their boundaries.

Recall that a *sheet* is defined as the closure (including semifree placements) of a connected component of $M(s_1,s_2)$ for some separated objects s_1 and s_2, and furthermore that without the assumptions of general position that there might be placements in $Vor_d(\Omega)$ not on any sheet. In this section we show that this anomaly will be excluded by the assumption of general position. We shall also characterize the interior and boundary of a sheet. We shall begin by classifying the interior and boundary of $M(s_1,s_2)$ relative to $Vor_d(\Omega)$. This presents some difficulties at placements where the ladder is near to a wall and an incident corner; therefore such placements need to be considered separately.

Definition. Let s_2 and s_3 be, respectively, a corner and an incident wall. A placement Z in *SFP* is *marginal* with respect to s_2 and s_3 if either (i) $R(Z)$ intersects both s_2 and s_3 or (ii) Z is free, $R(Z)$ touches s_2, and the line joining s_2 to the closest point on the ladder is normal to s_3, or (iii) Z is nonfree, one end of the ladder touches s_2, and the angle between the ladder and s_3 (measured from inside Ω) is $\leq \pi$.

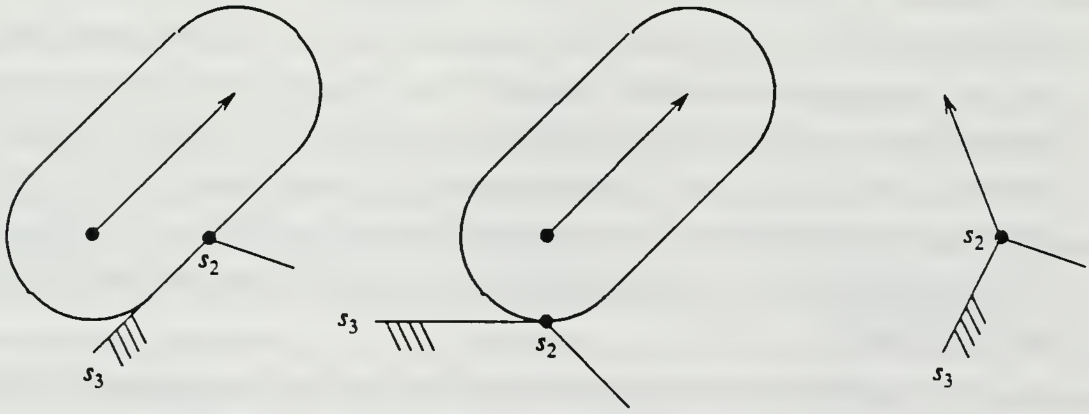

Case (ii)

Case (iii)

Case (i)

Figure 11. Marginal placements

When a marginal placement Z is in $Vor_d(\Omega)$ (so $Near(Z)$ intersects separated objects s_1 and s_2 but is marginal, say, with respect to s_2 and an incident wall s_3), then we shall see that Z is a placement common to two or more sheets. It is not difficult to show that the set of marginal placements is closed in FP (one merely shows that its complement is open). By definition, if s_1 and s_2 are both walls then $M(s_1,s_2)$ cannot contain any marginal placement, but they can exist in the other cases.

Lemma 4.1. Given two separated objects s_1 and s_2, the set of all non-marginal placements in $M(s_1,s_2)$ is locally path-connected and open in $Vor_d(\Omega)$.

Proof. Let Z_0 be a non-marginal placement in $M(s_1,s_2)$. It is enough to show that there exists an open neighborhood U of Z_0 in FP such that every placement in U is non-marginal, and that $U \cap Vor_d(\Omega)$ is a path-connected subset of $M(s_1,s_2)$. Let R be a thickening of $R(Z_0)$; thus $R \cap \Omega^c$ has two convex components C_i each intersecting s_i ($i = 1,2$). Let W be a neighborhood of Z_0 with the property that for every placement Z in W (i) Z is free; (ii) $R(Z) \subseteq R$ (corollary to Lemma 2.1); (iii) Z is not marginal (this imposes achievable constraints on the orientation of Z and the position of $P(Z)$); and (iv) the set of points in C_i closest to $B(Z)$ is contained in s_i (this is achievable since either s_i is a wall containing $C_i \cap \partial\Omega$ or s_i is a corner, closest to a unique point q on $B(Z)$, and the angles formed by qs_i with the incident walls are both greater than $1/2\pi$). Clearly, it follows that W intersects $Vor_d(\Omega)$ only in non-marginal placements in $M(s_1,s_2)$. Hence it only remains to find an open subset U of W such that U contains Z_0 and $U \cap Vor_d(\Omega)$ is path-connected. Note that $W \cap Vor_d(\Omega) \subseteq M(s_1,s_2) \subseteq VM(s_1,s_2)$ so it is enough to ensure that $U \cap VM(s_1,s_2)$ is (path)-connected. Consider neighborhoods U of Z_0 having the form $(x_1,x_2) \times (y_1,y_2) \times (\theta_1,\theta_2)$ and contained in W. Referring to Lemma 3.2, if we can ensure that either the (x,θ)- or the (x,y)-parametrization is defined on U then we shall be finished, since such a parametrization establishes a homeomorphism of $U \cap Vor_d(\Omega)$ with a set which is obviously path-connected. Clearly, if Z_0 is not a broadside placement then we can assume the same holds for all placements in U and hence the (x,θ)-parametrization is valid for $U \cap VM(s_1,s_2)$. Finally, when Z_0 is a broadside placement the (x,y)-parametrization is valid if U is sufficiently small (Lemma 3.2, corollary). Q.E.D.

Next we demonstrate the (seemingly obvious) fact that if in a free placement Z the racetrack touches more than two separated objects then it cannot be interior to any sheet.

Lemma 4.2. Let Z be a free non-marginal placement in which the racetrack intersects three or four objects s_i. Let (s_1,s_2) be any pair of these three or four objects where s_1 and s_2 are adjacent in cyclic order around the racetrack. Then

(i) every neighborhood U of Z intersects the interior of $M(s_1,s_2)$;

(ii) Z is in an (s_1,s_2)-sheet;

(iii) Z is not in the interior of any sheet.

Proof. Fix an open neighborhood U of Z in FP. Without loss of generality we may assume that there are four objects touching the racetrack, and that (with respect to the standard coordinate system for $VM(s_1,s_2)$) the other two objects touch the eastern segment of $\partial R(Z) \sim (s_1 \cup s_2)$, in the notation of Lemma 3.4 (the local monotonicity lemma). Thus the western segment does not touch any object, and it is clear that this holds true for all placements west of Z and sufficiently close to it on the same virtual streamline in $VM(s_1,s_2)$. Thus there exists a path $\{Z_t : t \in [a,b]\}$, where $Z = Z_b$, consisting of non-marginal placements west of Z in this virtual streamline and in which the western segment of $\partial R(Z_t)$ does not intersect $\partial\Omega$. We may assume that this path of placements is inside the open set U. By Lemma 3.4, if the interval $[a,b]$ is sufficiently small, then either (a) for all t in $[a,b)$ the eastern segment of $\partial R(Z_t)$ does not touch $\partial\Omega$ or (b) one of the two obstacles under consideration, s_1 say, is a corner, the placement Z is a broadside placement with respect to s_1 and some corner s_3 distinct from s_2, and $Near(Z_t)$ intersects s_3 for all t in $[a,b]$ (note that, assuming the interval $[a, b]$ is sufficiently small, the assumption of general position eliminates the possibility that $R(Z_t)$ could intersect all four objects when $a \le t < b$).

In case (a) every placement Z_t is in $M(s_1,s_2)$, and indeed the path Z_t consists of placements all in the same path component of the interior of $M(s_1,s_2)$: in this case, (i) and (ii) follow, since Z must be in the closure of the interior of a connected component of $M(s_1,s_2)$.

In case (b) we must be more careful. Observe first that Z cannot be a broadside placement with respect to s_1 and s_2. Let R be a thickening of $R(Z)$ and let C be the component of $R \cap \Omega^c$ intersecting the fourth object that intersects $Near(Z)$. Without loss of generality, a may be fixed so that $R(Z_t)$ is interior to R for each t in $[a,b]$, and only touches C for $t=b$. Fix any b' in (a,b). Since $R(Z)$ is continuous in Z, for every t in $[a,b']$ there exists an open neighborhood V of Z_t in $VM(s_1,s_2)$ of the form $\Phi((x_0,x_1)\times(\theta_0,\theta_1))$, where Φ is the canonical (x,θ)-parametrization of $VM(s_1,s_2)$, such that V is contained in U and $R(Z') \subseteq R \sim C$ for every placement Z' in V. Furthermore, within these sets it is easy to show that a placement $\Phi(x,\theta)$ is in $M(s_1,s_2)$ if and only if $\theta < \theta(Z)$. Since $[a,b']$ is compact, we can cover it by a finite set of such open sets, and hence (by choosing the maxima and minima, respectively, of the different bounds θ_0 and θ_1 used to define sets in this finite cover), we deduce that there exists a set of the form $\Phi([a,b']\times(\theta'_0,\theta'_1))$ contained in U and such that $A(b') = \Phi((a,b')\times(\theta'_0,\theta(Z)))$ is in $M(s_1,s_2)$. Clearly, $A(b')$ is an open path-connected set of placements in $M(s_1,s_2)$, and it intersects U. Thus, (i) follows. To show (ii), observe that all the sets $A(b')$ intersect pairwise, so they belong to the same path-connected component of $M(s_1,s_2)$; since Z is the limit of such placements $Z_{b'}$, and each $Z_{b'}$ is in the closure of

$A(b')$, it follows that Z is in an (s_1,s_2)-sheet. This concludes the proof of (i) and (ii).

Finally, (iii) is relatively trivial. Every neighborhood U of Z intersects the interiors of four different midsurfaces. No placement interior to an (s_1,s_2)-midsurface can be in any sheet other than an (s_1,s_2)-sheet. Hence every neighborhood of Z intersects at least four distinct sheets, so Z cannot be interior to a sheet. Q.E.D.

Lemma 4.3. Let Z be a marginal free placement in an (s_1,s_2)-sheet K. Then Z is not in the interior of K.

Proof. By definition, the placement Z cannot be a broadside placement: if a broadside placement were marginal then a wall incident to s_2, say, would be parallel to the line joining s_1 and s_2, contradicting the principle of general position. Therefore by Lemma 1.1, $P(Z) \in Vor_0(\Omega_\theta)$ where $\theta = \theta(Z)$.

We may suppose without loss of generality that the placement is marginal with respect to a corner s_2 and an incident wall s_3. (In what follows the proof will hold good for s_2 and s_3 equally.) There are two cases to consider: (i) where $Near(Z)$ intersects s_2 but not s_3, and (ii) where $Near(Z)$ intersects both s_2 and s_3.

(i) Suppose without loss of generality that the shortest line joining $B(Z)$ to s_2 is the line segment from $P(Z)$ to s_2 and this segment is normal to a wall s_3 incident to s_2. Note that s_2 and s_3 form part of $\partial\Omega_\theta$. There exist objects s and t (which could be the same and equal to s_1) forming part of $\partial\Omega_\theta$ such that $P(Z)$ is equidistant from s, t, s_2, and s_3 (s and t can be different when Z is a Voronoi vertex; we have not assumed that $Near(Z)$ has exactly two connected components; either $R(Z)$ could touch a third separated object or the placement could also be marginal with respect to s_1). Here, $P(Z)$ is a vertex of the augmented standard Voronoi diagram (in the sense of Kirkpatrick, see section 1) representing a point in $Vor_0(\Omega_\theta)$ common to the closures of the Voronoi 'cells' $V(s_2)$ and $V(s_3)$. We may assume that it is incident to an (s,s_2)-edge e and a (t,s_3)-edge f of $Vor_0(\Omega_\theta)$ (in the sense of Kirkpatrick) Again, there exists an object s' in S determining s in the sense that s is a subset of s' or $s' - (dCos\theta, dSin\theta)$. Similarly, there is an t' in S determining t.

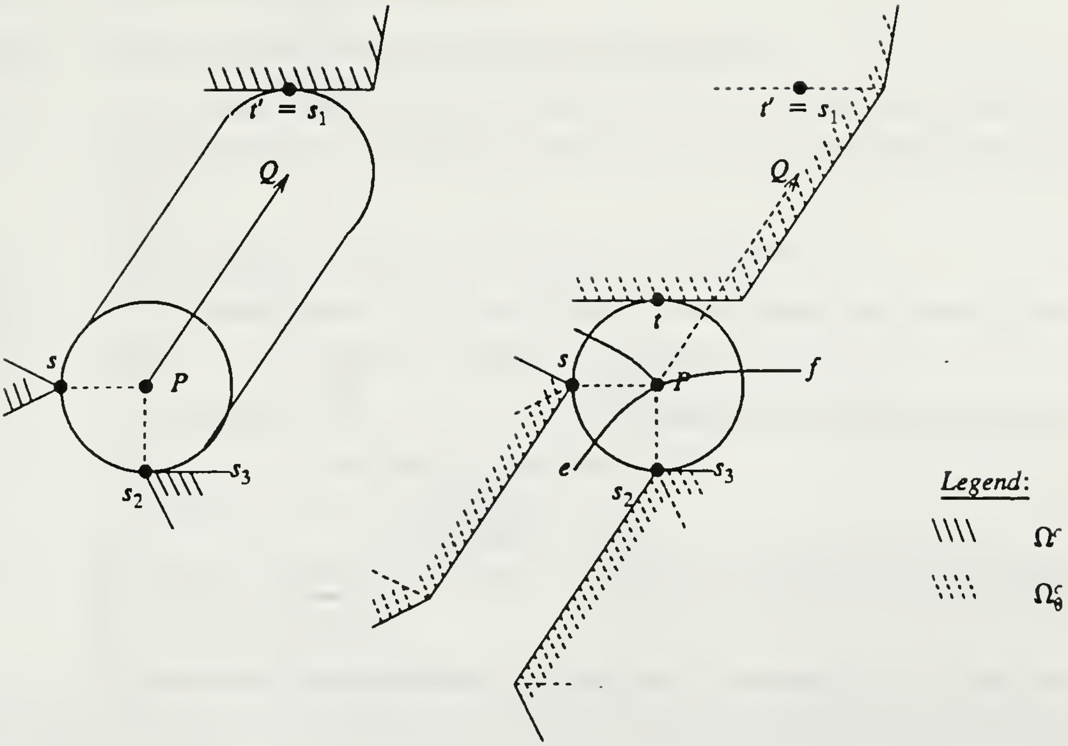

Figure 12. Illustrating case (i) with $t' = s_1$, say

It is easy to see that for points p close to $P(Z)$ in e, (p,θ) is not marginal and is interior to $M(s_2,s)$, and for points q close to $P(Z)$ in f, (q,θ) is not marginal and is interior to $M(s_3,t)$. Hence, as in the previous lemma, the placement Z is in the common boundary of two (or more) sheets and not interior to any sheet.

(ii) Here the treatment is slightly different. By the assumption of general position, the placement cannot be a broadside placement with respect to any of the objects touching $R(Z)$. Let s_2 be a corner and suppose that the wall s_3 incident at s_2 intersects the lower side of $R(Z)$ to the east of s_2 (with respect to the standard parametrization of $VM(s_1,s_2)$). Let s be the first object separated from s_2 at $\partial R(Z)$ as one moves clockwise around the boundary of $R(Z)$ from s_2; similarly t is the first object (the first object in this case is necessarily separated) at $\partial R(Z)$ anticlockwise around the boundary of $R(Z)$ from s_2. Note that possibly s and t are equal to s_1.

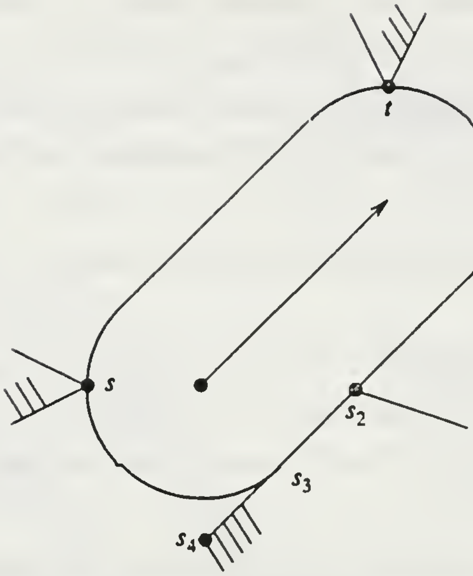

Figure 13. Illustrating case (ii)

There are two subcases here: (*a*) where the other end s_4 of s_3 is closest to some point in the interior of the ladder, and (*b*) where it isn't. Case (*a*) is essentially the same as when s_2 and s_4 are separated, so the placement is broadside with respect to these corners, and we may regard it as having been covered in case (b) of Lemma 4.2. By the same reasoning, in case (*b*) we can tell that Z is in the boundary of an (s, s_3)-sheet and also in the boundary of an (t, s_2)-sheet, and hence cannot be interior to any sheet. Q.E.D.

Let s_1 and s_2 be two separated objects, and K be an (s_1,s_2)-sheet. In general, a third object s_3 can determine part of the free boundary of K in a limited number of ways. In general, we define the *midline* $M(s_1,s_2,s_3)$ as the set of all placements Z in *SFP* such that either

(*a*) *Near*(Z) intersects all three objects (whether or not s_3 is separated from the other two objects), or

(*b*) *Near*(Z) intersects two of the objects of which one is a corner, $\{p\}$, say, the third object is a wall incident to p, and the shortest line joining p to the ladder is normal to that wall (ie., Z is marginal).

When Z is a nonfree placement, condition (*b*) is replaced by the requirement that one end of the ladder should touch p. The following list classifies the various kinds of midline. (In the classification, it is not assumed that the racetrack touches no other object. The notation has been chosen for mnemonic purposes: in particular, the primed letters indicate nonseparated objects of type W (wall) or C (corner).)

[WWW]	The Z-racetrack touches three walls.
[WWC]	The Z-racetrack touches two walls and a separated corner.
[WW′C′]	The Z-racetrack touches a wall and a corner separated from it, and the shortest line connecting the ladder to the corner is normal to one of the walls incident to that corner.
[WCC]	The Z-racetrack touches a wall and two corners, all separated.
[W′C′C]	The Z-racetrack touches two separated corners, and the shortest line connecting the ladder to one of the corners is normal to a wall incident to that corner.
[WC′C′]	The Z-racetrack touches a wall and two corners separated from the wall but not from one another (so it contains the wall connecting them).
[CCC]	The Z-racetrack touches three separated corners.
[CC′C′]	The Z-racetrack touches three corners, of which two are not separated (so it contains the wall connecting them).

It is a straightforward consequence of the 'thick walls' assumption that no semifree placement can be interior to any sheet; the proof is left to the reader. Bearing this in mind, the following theorem essentially summarizes the analysis of this section. It is a straightforward corollary of Lemmas 4.1, 4.2, and 4.3.

Theorem 4.4. The interior of every (s_1,s_2)-sheet K coincides with its intersection with the set of non-marginal placements in $M(s_1,s_2)$; distinct sheets can intersect only along their boundaries; the free boundary of every sheet consist of placements in one of the above eight categories. ●

In every placement Z in $Vor_d(\Omega)$, $R(Z)$ intersects at most four distinct objects. All of these cases have been covered in this section (relying somewhat on the assumption of general position and thick walls), and therefore the following result holds.

Theorem 4.5. $Vor_d(\Omega)$ is contained in the union of its sheets. ●

Remark. It is not difficult to apply the methods of this section to show that every Voronoi vertex is common to a bounded number of sheets. By the usual considerations of algebraic geometry, for any four objects s_i (in general position), there are a bounded number of placements Z in which the ladder is equidistant from them. Hence there are $O(n^4)$ Voronoi vertices and sheets, where $n = |S|$. (A more careful analysis shows the sharper estimate $\Omega(n^2)$ and $O(n^2 \log^* n)$; details are given in [OSY2]).

5. Connectivity of streamlines.

Recall that a streamline C in a (s_1, s_2)-sheet K is defined as the intersection of K with a virtual streamline X. The virtual streamlines partition $VM(s_1, s_2)$ into disjoint sets homeomorphic to closed unbounded intervals of real numbers. Since K is compact, so is C. In this section we show that (i) C must be homeomorphic to a closed interval $[a, b]$, (ii) the intersection C° of C with the relative interior K° of K is empty or homeomorphic to an open interval (a', b'), and (iii) $a = a'$ and $b = b'$ for all but finitely many 'degenerate' streamlines. We call $[a, a']$ and $[b', b]$ the *east* and *west ends* of the streamline respectively (in what follows, the 'easterly' direction means the direction of increasing x relative to the standard parametrization of $VM(s_1, s_2)$.) If C° is empty then we define both the east and the west end of C to be equal to C itself.

The simple principle underlying our somewhat tedious proofs is summarized as follows:

Lemma 5.1. Let C and its relative interior C° be as discussed above.

(i) If for any placement Z in $C \sim C^\circ$, either C° contains no placements west of Z or C° contains no placements east of Z, then C° is connected.

(ii) If for every (closed) subinterval $\{Z_t : t \in [a, b]\}$ of C, $M(s_1, s_2)$ contains no placements east of Z_b, then C is connected. ●

Lemma 5.2. Every y-streamline in a sheet K is a connected interval of (broadside) placements which intersects ∂K only at its endpoints.

Proof. Let Z lie in a (s_1, s_2)-sheet K. Since it is a broadside placement with respect to the corners s_i, it cannot be a marginal placement (defined in Section 4) with respect to those corners (this would contradict the principle of general position). First suppose that Z is a placement that is westernmost in its interval of $K \cap X$ where X is the virtual (s_1, s_2)-streamline containing Z. It is easy to see that if Z is free than Z is near a third object s. Furthermore, for all Z' on the streamline X west (resp. east) of Z, Z' is closer to (resp. farther away from) s than to (resp. from) s_i. The result then follows from Lemma 5.1. If Z is not free then the ladder touches two corners, which by convention we suppose are located on the y-axis. Since the two corners are not at distance d apart (general position), at least one of these corners touches the interior of the ladder, and so the walls incident to that corner are on the same side of the y-axis. Suppose that they are to the left of the y-axis. Then for every placement west of Z on the virtual y-streamline, the virtual Z-racetrack intersects the interior of Ω^c, and again no such placement can be in K. Q.E.D.

The situation is somewhat more complex when we consider θ-streamlines for a corner-corner sheet. This time there are no nonfree placements, but there is a possibility of degeneracy in that streamlines may intersect the boundary in a continuum of placements. However, the east and west ends are still connected, and only finitely many streamlines are degenerate in this way.

Lemma 5.3. Let K be an (s_1,s_2)-sheet where s_1 and s_2 are corners. Then every θ-streamline in K is connected and the intersection of every θ-streamline with the interior of K is connected. With finitely many exceptions, ∂K intersects the θ-streamline only at the two extremal placements of the streamline.

Proof. Let Z_0 be a placement in $X \cap \partial K$ where X is a virtual θ-streamline; note that Z_0 must be a free placement. Recall that according to the analysis in Section 4 either (i) a third point p in Ω^C touches the Z_0-racetrack or (ii) Z_0 is marginal in K.

In case (i) suppose w.l.o.g. that p touches $\Gamma(Z_0)$. Hence for all placements Z' east of Z_0 in X, we have p in $\Gamma(Z')$ and hence in the virtual Z_0-racetrack. This shows that no placement east of Z_0 in X can be interior to K (nor to any other component of $M(s_1, s_2)$). Furthermore, unless p is in a corner touching the same side of the racetrack as one of the other corners, no placement east of Z_0 in X can be in the closure of $M(s_1, s_2)$, hence no such placement can be in K. Thus Z_0 is the eastern endpoint of $\partial K \cap X$ except for at most finitely many values of θ, and the set of placements Z' east of Z_0 in X does not intersect the interior of K.

In case (ii) a wall s_3 incident to s_2 is tangent to the Z_0-racetrack. First suppose that s_3 is not parallel to the ladder, so s_1 is closest to one of the ends of the ladder, Q, say. We may assume that s_3 extends eastwards from the y-axis. Then for all placements Z east of Z_0 in X which are sufficiently close to Z_0, $sep(s_1, B(Z)) = sep(s_1, Q(Z))$, but the angle $s_3 s_1 Q(Z)$ is acute, so s_3 intersects the interior of $\Gamma(Z)$. Thus by the monotonicity lemma s_3 intersects the interior of $\Gamma(Z)$ for every Z east of Z_0 in X and Z_0 is an eastern endpoint of the interval of the streamline containing it. Finally, Z_0 may be a marginal placement in which the ladder is parallel to s_3, so s_3 lies in the tangent to the side of the racetrack. Assuming again that s_3 extends eastwards from the y-axis, we know that for all placements Z in X east of Z_0 and sufficiently close to it, $\Gamma(Z)$ intersects s_3. Hence this holds for all placements Z east of Z_0, so this set of placements cannot intersect the interior of K.

Thus we have shown that for the cases considered all placements Z in $X \cap K^\circ$ lie to the west of Z_0; the other cases are covered by symmetry, ie., where in case (i) p touches $\Sigma(Z_0)$ or in case (ii) s_3 extends westwards from the y-axis, so Lemma 5.1 implies that $X \cap K^\circ$ is connected. Furthermore, there are at most finitely many degenerate streamlines. Suppose that Z_0 is in a degenerate streamline, with $K^\circ \cap X$ lying to the west of Z_0, and suppose Z_1 is the eastern endpoint of that interval of $K \cap X$ containing Z_0. Then the previous discussion shows that either case (i) or case (ii) obtains; in case (i) the point p' under consideration cannot touch a side of $R(Z_1)$ if that side touches s_1 or s_2 (by general position) and in case (ii) either the ladder is not parallel to the wall s_3 under consideration, or it it, and the closest point to s_1 on $B(Z_1)$ is an endpoint of the ladder. Following the same reasoning as before, no placement east of Z_1 in X can be in K, so the entire result follows from Lemma 5.1. Q.E.D.

The next task is to prove a similar result about wall-wall sheets. In this case (and in the wall/corner case) we do not have the fact that $R(Z) = \Gamma(Z) \cup \Sigma(Z)$, so the argument becomes more delicate at one point. The following rather simple lemma compensates for the difficulty.

Lemma 5.4. Suppose that K is an (s_1,s_2)-sheet, where both objects are walls, and p is a point interior to the complement Ω^c of Ω. Then if $p \in \Gamma(Z_0)$ (respectively, $\Sigma(Z_0)$) for *some* Z_0 in K then $p \in \Gamma(Z)$ (respectively, $\Sigma(Z)$) for *every* Z in K.

Proof. We shall only prove this for the sets $\Gamma(Z)$; the other part of the lemma is proved in exactly the same way. We shall assume that the result is not true and obtain a contradiction: namely, that if Z_0 and Z_1 are placements in $VM(s_1,s_2)$, where $p \notin \Gamma(Z_1)$, and $\{Z_t : t \in [0, 1]\}$ is a continuous path in $VM(s_1,s_2)$ connecting these placements, then for some placement Z' along this path the Z'-racetrack intersects the interior of Ω^c.

The point p must of course be interior to the wedge J discussed in Section 3. Choose any ϵ less than $d(p,\Gamma(Z_1))$, and such that the closed disc D_ϵ of radius ϵ around p is interior to the complement Ω^c of Ω. Note (Lemma 3.5) that the distance $d(p,\Gamma(Z_t))$ varies continuously with t, and is zero for $t = 0$. Therefore there exists some placement Z' along the path for which $d(p,\Gamma(Z'))$ is nonzero but less than ϵ. This means that $p \notin \Gamma(Z')$ but $\Gamma(Z')$ intersects the interior of D_ϵ. It follows easily by convexity arguments that there exists some point common to the boundary of $\Gamma(Z')$ – interior to the wedge J – and the interior of D_ϵ. Clearly, such a point is in the virtual Z'-racetrack and hence Z' cannot be in the closure of $M(s_1,s_2)$. Q.E.D.

Note: the placement Z' may be 'virtually' semifree in $VM(s_1, s_2)$, in which case the ladder $B(Z')$ intersects the interior of Ω^c; Z_0 and Z_1 need not be free. This point will be of significance when we analyze the semifree boundary of a sheet (Section 7)

Lemma 5.5. Let K be an (s_1,s_2)-sheet where s_1 and s_2 are both walls. Then every streamline in K is connected and the intersection of every streamline with the interior of K is connected. ∂K intersects the streamline in two isolated placements at each end of the streamline, with the exception of at most four degenerate orientations.

Proof. Let C be a θ-streamline and let X be the virtual streamline containing it. We suppose that the objects s_1 and s_2 are arranged as stipulated in Section 3. First note that the streamline can contain at most one nonfree placement and clearly such a placement must be westernmost in C (indeed, it is westernmost in X). Next let us consider a free placement Z_0 in $X \cap \partial K$ such that

$Near(Z_0)$ has two connected components: this corresponds to the case [WWC′] discussed in Section 4. Suppose for definiteness that $Near(Z_0)$ intersects the east endpoint s_3 of s_2. Clearly, for all placements Z east of Z_0 in X, $Near(Z)$ cannot intersect s_2 without also intersecting s_3, so no placement east of Z_0 is interior to K. Note that Z_0 is unique in this case unless the ladder is parallel to s_2 (which happens at four orientations). (This essentially covers the case [WWC′].)

The last possibility is that $R(Z_0)$ intersects another point q of $\partial\Omega$ in the interior of the wedge J. We may suppose that q touches $\Gamma(Z_0)$ (the case for Σ is the same). In this case no placement Z east of Z_0 in X can be in K: for suppose that Z_1 is a placement east of Z_0 in the closure of $M(s_1,s_2)$. It is clear, by the thick walls principle, that there is a point p (close to q) in the interior of Ω^c such that p is in $\Gamma(Z_1)$ but not in $\Gamma(Z_0)$. Hence by Lemma 5.4, Z cannot be in K. Q.E.D.

Finally, we consider the case of a wall-corner sheet. The property stated in Lemma 5.4 holds in this case too, and the proof is the same, so we shall just state the result.

Lemma 5.6. If K is an (s_1,s_2)-sheet and p a point interior to the complement of Ω, where s_1 is a corner and s_2 a wall, and for some placement Z_0 in K, $p \in \Gamma(Z_0)$ (respectively, $\Sigma(Z_0)$), then for all placements Z in K, $p \in \Gamma(Z)$ (respectively, $\Sigma(Z)$).

Lemma 5.7. Let K be an (s_1,s_2)-sheet where s_1 is a corner and s_2 is a wall. Then every streamline in K is connected and the intersection of every streamline with the interior of K is connected. With finitely many exceptions, ∂K intersects the streamline in the two isolated placements at each end of the streamline.

Proof. Let us consider a placement Z_0 in $\partial K \cap X$, where X is a virtual (s_1,s_2)-streamline. Suppose first that Z_0 is in K and is not a free placement. Thus the ladder touches both objects. The simplest case to consider is where s_1 touches the ladder at a point interior to the ladder. Without loss of generality both walls incident to s_1 extend west of the line through the ladder (since the ladder cannot be horizontal, this makes sense). It is clear, therefore, that for any placement Z in X, west of Z_0 and sufficiently close to it, $B(Z)$ and hence $\Delta(Z)$ intersects Ω^c in its interior. Hence since $\Delta(Z)$ expands as Z moves westward, no placement whatsoever west of Z_0 in X can be in K. Next it is possible that s_1 touch an endpoint of the ladder. Without loss of generality, at least one of the walls s_3 incident to s_1 extends westward from the line through the ladder, and Q is the end of the ladder which touches s_1. As Z moves westward along X, Q moves away from s_1 along a straight line at angle $1/2(\pi+\theta)$ where θ is the orientation determining X. If the line through s_3 meets the line through s_1 at a point west of P then the ladder transfixes the wall when Z is close to Z_0 and the arguments are as for the simplest case. Otherwise Q moves closer to this wall than to s_1 and Q is locally the point on the ladder closest to this wall, and again the previous arguments involving $\Delta(Z)$ apply. This disposes of the case where Z_0 is a semifree placement.

Next suppose that $Z_0 \in M(s_1, s_2)$ but is a marginal placement. Without loss of generality, Z_0 is east of s_1. There are two possibilities: (a) some wall s_3 incident to s_1 is tangent to $\Delta(Z_0)$ and (b) Z_0 is a marginal placement with respect to s_2 and the eastern endpoint of s_2. First consider (a). It follows by the usual arguments that for placements Z in X close to Z_0 and east of the wall either remains tangent to (in which case the ladder is parallel to it) or penetrates the interior of $\Delta(Z)$. This means that no placement in X east of Z_0 and close to it is interior to K, and the growth property of $\Delta(Z)$ makes this true for all placements in X east of Z_0. Case (b) is easily disposed of as in the wall-wall case.

Therefore we have dealt with all cases where the ladder is in ∂K but $Near(Z_0)$ has only two components, and so we can now consider the case where some point p in $\partial\Omega$ touches the racetrack outside $s_1 \cup s_2$. Without loss of generality either (a) the point p is above the line $y=h$ and the ladder is east of s_1 or (b) the point p touches the racetrack on the boundary of $\Gamma(Z_0)$. In either case we determine that no placement east of Z_0 in X can be interior to K, and indeed no placement east of Z_0 can be in K unless $\{p\}$ and s_1 both touch the same side of the racetrack. In this latter case, we have that the component of $\partial K \cap X$ that contains Z_0 is a continuum of placements. Thenceforth the arguments are as for the previous lemmas. Q.E.D.

We may summarize the results of this section in the following general proposition, which is only a restatement of the relevant lemmas.

Proposition 5.8. In any sheet K, every streamline is connected and the intersection of every streamline with the interior of K is connected. The intersection of every streamline with ∂K consists of two closed intervals, called the *east* and *west ends* of the streamline, and these intervals coincide if and only if the streamline does not intersect the interior of the sheet. With finitely many exceptions, each end of the streamline is a singleton set.

6. The ends of a sheet.

In this section we shall define the ends of a sheet and demonstrate their connectivity. From there it is not difficult to define the semifree skeleton of a sheet and prove that it is connected.

Definition. The *east* (*west*) *end* of a sheet K is the union of the east (respectively, west) ends of all the streamlines in K. From the results of the previous section, the boundary ∂K of K is the union of its ends. An orientation θ is *admissible* (for K) if there exists some placement in K at orientation θ. In general, this means that the θ-streamline in K is nonempty; for broadside placements it means that at least one of the y-streamlines for that orientation is nonempty. $\Theta(K)$ denotes the set of orientations admissible for K. Being the continuous image of a compact connected set of placements, it is a compact connected set of angles, i.e., either the whole range of S^1 (the unit circle) or a closed interval $[\theta_{min},\theta_{max}]$.

Notice that by the analysis of the preceding section, for only finitely many orientations θ in $\Theta(K)$ do the θ-streamlines intersect the east or west end of the sheet in more than one placement – in which case that intersection is a connected continuum of placements: let us call such orientations *degenerate* (we exclude broadside orientations, for which such degeneracy is impossible). We use $E(\theta)$, $E(y)$, (respectively, $W(\theta)$, $W(y)$) to denote the easternmost (respectively, westernmost) placement in the θ-streamline (or y-streamline). Again, we use $EE(\theta)$ ($WE(\theta)$) to denote the east and west ends of the θ-streamline. Thus $EE(\theta) = \{E(\theta)\}$, and $WE(\theta) = \{W(\theta)\}$ for all but finitely many (non-broadside) orientations θ in $\Theta(K)$. Also write

$$EE(\Phi) = \bigcup \{ EE(\phi) : \phi \in \Phi \}$$

for any $\Phi \subseteq \Theta(K)$, and define $WE(\Phi)$ similarly.

Lemma 6.1. Let I be a subset of $\Theta(K)$ not containing any broadside orientations nor any degenerate orientations. Then the restrictions of the functions $E(\theta)$ and $W(\theta)$ to I are similarly homeomorphisms.

Proof. It is enough to prove the result for the map E, since the treatment of the other case is identical. Let θ_n converge to θ_0 in I. By passing to a subsequence if necessary, we can assume that the placements $E(\theta_n)$ converge to a placement Z_0 in SFP. Clearly, Z_0 is a semifree placement common to the boundary of K and the θ_0-streamline in K. Since the θ_0 is nondegenerate and not a broadside placement, it intersects ∂K in one or two placements: it intersects ∂K in one placement only when the streamline consists of one Voronoi vertex. Hence, if $Z_0 \neq E(\theta_0)$ then $Z_0 = W(\theta_0)$; the θ_0-streamline intersects the interior of K, so there is an open set V in K, of the form $\Phi((x_1,x_2)\times(\phi_1,\phi_2))$, say (where Φ is the standard (x,θ)-parametrization of the sheet), such that V intersects the θ_0-streamline, so the x-coordinate of Z_0 is x_1 or less. However, for sufficiently large n, $E(\theta_n)$ have x-coordinate x_2 or more, which is absurd. Thus the maps E and W are continuous on I. Since the projection mapping $Z \rightarrow \theta(Z)$ commutes with both E and W, the restrictions of the maps to I are homeomorphisms. ●

In a corner-corner sheet K the set of broadside placements is defined by one or two orientations (depending on the range of $\Theta(K)$). Thus there are one or two components in the set, which we shall denote J^+ and J^- (corresponding to broadside orientations $\pm\pi/2$). Without loss of generality we may concentrate on J^+. It is elementary to show that J^+ and J^- are path connected. Let $Y^+(K)$ be the range of values of y such that the y-streamlines in J^+ are nonempty. Since it is the continuous image of a compact connected set, it is a closed interval of real numbers. No y-streamline can intersect ∂K in more than two placements (Lemma 5.1), so the following lemma may be proved in exactly the same way as Lemma 6.1.

Lemma 6.2. Given J^+ and $Y^+(K)$ as above, the maps E and W are homeomorphisms between $Y^+(K)$ and the intersection of J^+ with the east (respectively, west) end of K. ●

Let us now investigate the transition between broadside and non-broadside placements in a corner-corner sheet K. Clearly $Y^+(K)$ is a subinterval of $[h-d, -h]$ where h is the distance of s_1 and s_2 from the origin of the standard coordinate system. Suppose Z_0 is a *free* broadside placement at the east end of the $(-h)$-streamline. (So $Z_0 = (x_0, -h, \pi/2)$ where $x_0 > 0$, since were $x_0 \leq 0$ it would equal zero and since Z_0 is easternmost, Z_0 would not be free.) Suppose there exists an orientation $\theta_0 < \pi/2$ such that $[\theta_0, \pi/2] \subseteq \Theta(K)$. Without loss of generality this range does not include any degenerate orientations, so the map E is a homeomorphism when restricted to $[\theta_0, \pi/2)$ (Lemma 6.1). Let Z be any placement interior to the sheet K and in the same streamline as Z_0. Appealing to the corollary to Lemma 3.2, there exists a neighborhood U of Z in $Vor_d(\Omega)$ on which the (x,y)-parametrization is valid; we can assume that $U = \Psi((x_1,x_2)\times(y_1,y_2))$ where Ψ is the (x,y)-parametrization. Clearly, $\Psi((x_1,x_2)\times(-h,y_2))$ intersects all θ-streamlines for θ in some small range $(\theta_1, \pi/2)$. Suppose that θ_n is a sequence converging to $\pi/2$ from below, and let Y_0 be any limit point of $E(\theta_n)$. By definition, for all sufficiently large n, $E(\theta_n)$ have x-coordinate $\geq x_2$. One can show also by simple geometric considerations (see, for instance, the development in appendix A), that the y-coordinates of placements are nondecreasing along a θ-streamline if $0<\theta<\pi/2$.

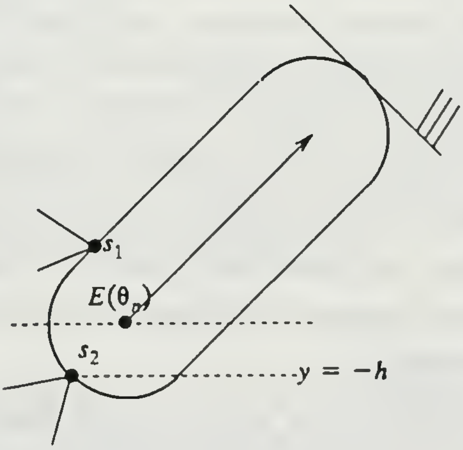

Figure 14.

This implies that $E(\theta_n)$ have y-coordinate $> -h$ for all sufficiently large n; hence Y_0 has y-coordinate $-h$ (which is maximal for the broadside placements at orientation $\pi/2$) and lies on the same y-streamline as Z_0. Since it has x-coordinate $\geq x_2$ and is in ∂K we conclude that it coincides with Z_0. Thus Z_0 is the unique limiting placement for $E(\pi/2^-)$. Similarly if the $(h-d)$-streamline contains placements $(x, h-d, \pi/2)$ with $x>0$ (implying it is nonempty), then $E(h-d)$ is the unique limiting placement of $E(\pi/2^+)$. Similar results apply to the western ends of the sheet, for which the requirement is that the sheet admit free broadside placements with the ladder in the negative half-plane $\{x<0\}$. Let us summarize as follows.

Lemma 6.3. In a corner-corner sheet K, let $[\theta_0, \theta_1]$ be a subinterval of $\Theta(K)$ containing a broadside orientation but no degenerate orientation. If the east end of the sheet contains only free placements, then the set of placements in the east end whose orientation lies in the given range is homeomorphic to a closed interval. The same holds for the west end of the sheet if it contains only free placements.

Proof. For simplicity let us assume that the given range only includes one broadside orientation, $\pi/2$, and this coincides with neither θ_0 nor θ_1 (it is trivial to enumerate all other possibilities). Since $[\theta_1, \frac{\pi}{2})$ is admissible the continuity of the map E implies $E(\frac{\pi}{2}^-)$ exists and hence the $(-h)$-streamline is non-empty. Since the lemma assumes $E(-h)$ is free, the preceding argument shows that $E(\frac{\pi}{2}^-) = E(-h)$. Similarly, $(\frac{\pi}{2}, \theta_2]$ is admissible implies $E(\frac{\pi}{2}^+)$ exists and equals $E(h-d)$. This in turn implies that $[-h, h-d] \subseteq Y^+(K)$ and hence equals it. Then we may write the east end of the sheet in the form $E([\theta_0, \pi/2)) \cup E([h-d, -h]) \cup E((\pi/2, \theta_1])$, or $e_1 \cup e_2 \cup e_3$ for short. Each of these sets is the homeomorphic image of the corresponding half-open or closed interval, they are disjoint, $\bar{e}_1 = e_1 \cup E(-h)$, and $\bar{e}_3 = e_3 \cup E(h-d)$; the result follows immediately. Q.E.D.

On the other hand, if Z_0 is a semifree placement where an end of the ladder touches one of the corners, then one can proceed by case analysis: there are essentially three cases to consider (see the figure below).

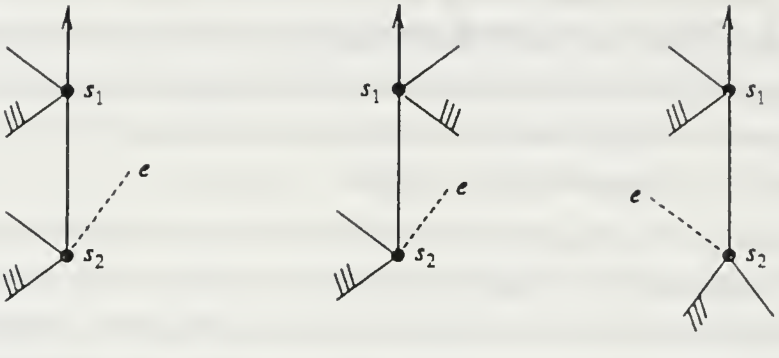

Case(a) Case(b) Case(c)

Figure 15. Illustrating limiting nonfree broadside placements

In case (a), all four walls incident to the two corners extend to the same side (westwards, say) of the y-axis; in case (b), the pairs of incident walls extend in opposite directions from the y-axis; and in case (c) three of the four walls extend in the same direction from the y-axis. Any other possibility (except what would result from reflecting Ω in the y-axis) is ruled out by the assumptions of general position and thick walls. Let s_3 be the wall incident to s_2 extending westward from s_2 and subtending the smaller angle with the positive y-axis, and let e be the normal to s_3 through s_2 and directed into Ω. For the three given paradigms this defines e uniquely. It is a straightforward exercise in geometry to show that for all points p in a sufficiently small initial segment of e there exists a unique placement $Z(p)$ in the (s_1,s_2)-sheet K such that $P(Z(p)) = p$. (Compare this with the corollary to Lemma 3.2.) In cases (a) and (c) these placements $Z(p)$ have orientation less than $\pi/2$ and in case (b) they have orientation greater than $\pi/2$. Moreover, it is easy to see that $Z(p)$ is always in the western end of the sheet: placements further west than $Z(p)$ in the same virtual streamline relative to s_1 and s_2 are strictly closer to s_3 than s_2. Thus Z_0 is the unique limit of $W(\theta_n)$ as θ_n converges to $\pi/2$ from above in case (b) and from below in the other two cases.

It is easy to show in case (c) above that Z_0 is also the limit of placements in the eastern end of K, $\pi/2$ is an upper bound to $\Theta(K)$, and Z_0 is the unique placement in the sheet at orientation $\pi/2$. Let $Z_1 = (0,h-d,\pi/2)$ with respect to the standard parametrization: it is the 'lowest' possible nonfree placement. By the same reasoning as before, in case (a) Z_1 is the unique limit of placements $W(\theta_n)$ as θ_n converges to $\pi/2$ from above; in case (b), $\pi/2$ is a lower bound to $\Theta(K)$ and Z_1 is the unique limit of placements of the form $E(\theta_n)$ as θ_n converges to $\pi/2$ from below. We shall have occasion to discuss these cases further in Section 7; for the present it is enough to note that the case-analysis allows us to extend Lemma 6.3 to cover the case of semifree limiting placements (one should note with some care that only in case (*a*) above is $\pi/2$ *not* a (lower or upper) bound for $\Theta(K)$, so the following lemma is carefully worded). Its proof is identical to that of Lemma 6.3.

Lemma 6.4. Let K be a corner-corner sheet and suppose that $[\theta_0,\theta_1]$ is a subset of $\Theta(K)$, not the full circle of angles, containing a broadside orientation but no degenerate orientation. Then the set of placements whose orientation is in the given range and which lie in the east (respectively, west) end of K is homeomorphic to a closed interval. •

We shall now consider the behavior of the functions E and W near their discontinuities (note that from the analysis of Section 5 there can be only finitely many such discontinuities). First define an *extremal streamline* in a sheet K to be either (i) a θ-streamline where θ is one of the bounds of $\Theta(K)$ (implying that not all orientations are admissible in K, and also that θ is not a broadside placement in the corner-corner case), or (ii) a y-streamline where y is one of the bounds of the set $Y^+(K) \cup Y^-(K)$ parametrizing the broadside placements but is not equal to $-h$ or $h-d$. It is very easy to see that an extremal y-streamline consists of exactly one nonfree placement Z in which the ladder touches two corners and is 'blocked' by having a wall touch one of its ends (see the proof of Lemma 5.2).

Lemma 6.5. If X is a streamline extremal in K then $X \subseteq \partial K$.

Proof. Clearly we need only show that X is not extremal in K if it intersects the interior of K. Let Z be a placement interior to K so by Theorem 4.4 it is free and non-marginal. By Lemmas 3.1 and 3.2 and their corollaries it lies in a set of the form $\Phi((x_1,x_2)\times(\theta_1,\theta_2))$ or of the form $\Psi((x_1,x_2)\times(y_1,y_2))$ where Φ and Ψ are the appropriate parametrizations, and the given (open) set lies entirely in K. If X is a θ-streamline then $\theta \in (\theta_1,\theta_2)$ whereas if it is a y-streamline then y is interior to (y_1,y_2); therefore X cannot be an extremal streamline. Q.E.D.

Remark: The converse to Lemma 6.5 holds also, as will become evident from the case analysis in Lemma 6.7 next. It should be emphasized that the above lemma shows that the streamline at an extremal orientation coincides with its east, as well as west, ends.

Let us now consider a degenerate orientation θ at which the θ-streamline intersects ∂K in two intervals of which at least one contains an infinite continuum of placements.

Lemma 6.6. Given an (s_1,s_2)-sheet K, let θ be an orientation in $\Theta(K)$ such that $EE(\theta)$ contains more than one placement, and suppose that it lies in a closed subinterval $[\theta_0,\theta_1]$ of $\Theta(K)$ which contains no other degenerate orientation. Then the set $EE([\theta_0,\theta_1])$ is homeomorphic to a closed interval of real numbers. Symmetric results hold when $WE(\theta)$ contains more than one placement.

Proof. Let $\{Z_t: a \le t \le b\}$ parametrize $EE(\theta)$. The analysis is best separated into several cases.

Case (a) [WCC, CCC] One of the objects, s_1, say, is a corner, and s_1 and another separated corner s_3 touch a side of the Z_t-racetrack (in the eastern segment according to the notation of Lemma 3.4) for all t in $[a,b]$. Let us first assume that s_1 and s_3 touch the western side of the racetrack. (This situation can happen in the corner-corner and wall-corner cases.)

In placement Z_a either (a.1) s_3 is closest to an end of the ladder, (a.2) Z_a is nonfree or another object s_4 touches the racetrack (necessarily in the western segment, by Lemma 3.4), or (a.3) Z_a is marginal with respect to s_2.

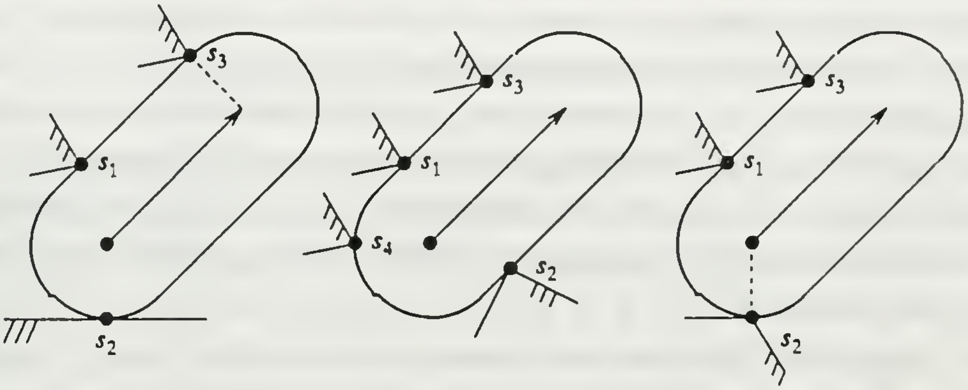

Case (a.1)

Case (a.2)

Case (a.3)

Figure 16. Illustrating placements Z_a

In case (a.1), Z_a is the unique limit of placements of the form $E(\theta_n)$ as θ_n converges to $\theta(Z_a)$ from above; this is a consequence of Lemma 6.3, if we observe that Z_a is a limiting broadside placement in $M(s_1, s_2, s_3)$ and the east end of the adjacent (s_3,s_1)-sheet. In cases (a.2) and (a.3), $Z_a = W(\theta)$ and θ is maximal, i.e., $\theta = \theta_1$: this follows from the analysis of Section 5, and it would be tedious to repeat the work.

Now let us consider the various ways in which Z_b can be extremal. Either (a.4) one end of the ladder is closest to s_1, or (a.5) the racetrack touches another object s_4, necessarily on its eastern segment, or (a.6) the placement Z_b is marginal with respect to s_2. In case (a.4) Z_b is easily seen to be the unique limit of placements of the form $E(\theta_n)$ as θ_n converges to θ from below. In case (a.4) the same holds, except that if Z_b is free then the placements $E(\theta_n)$ are simulataneously closest to s_1, s_2, and s_4. In case (a.2) the same holds, except that for all sufficiently large n the placements $E(\theta_n)$ are all marginal with respect to s_2.

Let us summarize these remarks. Assuming that s_1 and s_3 touch the western side of the Z_t-racetracks (with respect to the standard parametrization of $VM(s_1, s_2)$), then only if Z_a satisfies case (a.1) can $\theta \neq \theta_1$. There are three subcases. If $\theta = \theta_0$ then case (a.1) applies, $Z_a = W(\theta_1)$, $Z_b = E(\theta_1^-) = E(\theta_1)$, and

$$EE([\theta_0, \theta_1]) = EE(\theta_1) \cup E([\theta_0, \theta_1))$$

is homeomorphic to a closed interval as in Lemma 6.3. Again if $\theta_0 < \theta < \theta_1$ then case (a.1) applies, $Z_a = E(\theta^+)$, $Z_b = E(\theta^-)$ and again the result follows. Finally if case (a.1) does not apply then $\theta = \theta_1$, $Z_a = W(\theta_1)$ and $Z_b = E(\theta_1) = E(\theta_1^-)$ and again the result holds. This discussion will be complete when we cover the cases where s_1 and s_3 touch the eastern sides of the Z_t-racetracks. By our conventions of parametrization (bearing in mind what 'easterly' connotes), this is possible only when s_2 is a wall. In this case there are three differences: (1) In cases (a.2) and (a.5) Z_a cannot be semifree but Z_b can, and (2) $\theta = \theta_0$ is a lower limit except in case (a.1). The analysis continues as before, covering case (a).

Case (b) [WC'C', CC'C', WW'C', W'C'C] s_1 is a corner, and a wall s_3 incident to s_1 is tangent to a side of the racetrack. Again as in case (a), we first suppose s_1 to touch the western side of the racetrack. (Here $Near(Z_t)$ may also intersect the other endpoint s_4 of s_3 for some or all of the placements Z_t.) In placement Z_a either (b.1) s_1 is closest to an endpoint of the ladder, (b.2) Z_a is nonfree or the racetrack touches another object s_5, necessarily in its western segment, or (b.3) Z_a is marginal with respect to s_2.

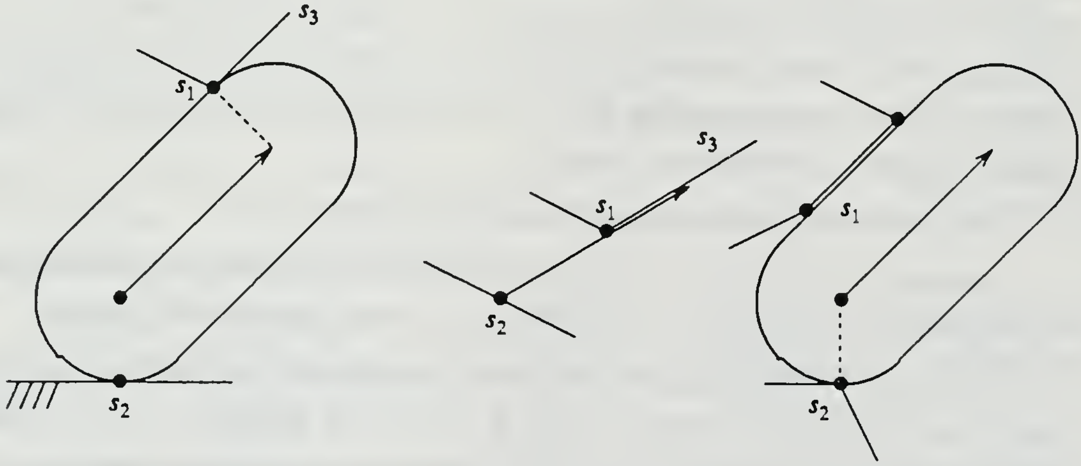

Figure 17. Illustrating placements Z_a in case (b)

In case (b.1) Z_a is the unique limiting placement $E(\theta^+)$; in the other two cases, as before, Z_a =

$W(\theta)$ and $\theta = \theta_1$. In placement Z_b either (b.4) s_1 is closest to an endpoint of the ladder, or (b.5) Z_b is nonfree or the racetrack touches another object s_5, necessarily in the eastern segment, or (b.6) the placement is marginal with respect to s_2. In any case, Z_b is the unique limit $E(\theta-)$. Thus the result holds as for case (a). The same remarks as given in case (a) when s_1 touches the eastern side may be verified, concluding case (b).

Case (c) [WW'C', WC'C']. Both objects are walls, and the placements Z_t are marginal with respect to one of them. Without loss of generality, $\theta = \phi$, the angle the wall s_1 makes with the positive x-axis in the standard coordinate system. Again let us consider the placement Z_a. In placement Z_a, either (c.1) the Q-end of the ladder is closest to the eastern corner of s_1, (c.2) the placement is semifree or the racetrack touches another object, necessarily in its western segment, or (c.3) the Q-end of the ladder is closest to the western corner of s_1 or the placement is marginal with respect to s_2. These cases are almost identical to those considered for case (b), and we deduce that in case (c.1) Z_a is the unique limit of $E(\theta^+)$ whereas in the other two cases the streamline is extremal. Similarly, in placement Z_b, either (c.4) the P-end of the ladder is closest to the eastern corner of s_1, (c.5) the racetrack touches another separated object, necessarily in the eastern segment, or (c.6) the placement is marginal with respect to s_2. In all cases, $Z_b = E(\theta^-)$, as before. This concludes the local analysis of the eastern end of K.

Except in the case where K is a wall-wall sheet, the standard coordinate system can be reflected in the y-axis, so the local analysis holds for the western end of K. In the wall-wall case we suppose again that $WE(\theta)$ is parametrized as an interval $\{Z_t\colon t\in[a,b]\}$. Cases (c.1-6) again describe the placements Z_a and Z_b. In cases (c.1-3), Z_a can be shown to equal $W(\theta^+)$. In case (c.4), Z_b can be shown to equal $W(\theta^-)$ and in the other two cases the streamline must be extremal. Hence the same reasoning applies to this case, and we deduce that the set $WE([\theta_0,\theta_1]$ is homeomorphic to a closed interval. Q.E.D.

We are now in a position to prove the main result of this section.

Lemma 6.7. The east (resp. west) end of a sheet K is a connected set homeomorphic either to a closed interval or to the unit circle S^1. The latter is true if and only if $\Theta(K) = S^1$.

Proof. It is sufficient to prove this for the east end of K. By the preceding lemmas we know that for any closed subinterval $[\theta_0,\theta_1]$ of $\Theta(K)$ the set $EE([\theta_0,\theta_1])$ is homeomorphic to a closed interval of real numbers. In particular, if $\Theta(K) \neq S^1$ then we can choose θ_0 and θ_1 to be the bounds of $\Theta(K)$. Otherwise we can choose an arbitrary pair θ_0 and θ_1 of non-degenerate and non-broadside orientations and separate $\Theta(K)$ into two subsets $[\theta_0,\theta_1]$ and $[\theta_1,\theta_0]$ (taken in counterclockwise order around S^1). These define subsets of the east end of the sheet, each of which is homeomorphic to a closed interval, and which intersect only at the endpoints; hence their union is clearly homeomorphic to S^1. Q.E.D.

The next lemma is an extremly simple consequence of the above analyses.

Lemma 6.8. If K admits an extremal streamline, then ∂K is path-connected.

Proof. An extremal streamline consists of placements common to both ends of the sheet (Lemma 6.5), so ∂K may be represented as the union of two intersecting path-connected sets. Q.E.D.

In the standard parametrization of a sheet K let us single out two canonical orientations: where the ladder is horizontal in the standard parametrization. We call the associated streamlines (which may, of course, be empty) the *core* streamlines. Finally define the *skeleton* of the sheet K to be the union of its ends with the core streamlines. Since each end of a sheet is connected and any streamline must intersect both ends, it follows at once that the skeleton of a sheet is connected.

Theorem 6.9. The skeleton of a sheet K is a connected subset of K, topologically homeomorphic to a network (ie., 1-dimensional complex), which contains ∂K. Since two sheets K_1 and K_2 can only intersect along their boundaries, it follows that the union of the skeletons of all sheets, which we call the *skeleton* of FP, is a network whose connectivity defines the connectivity of $Vor_d(\Omega)$ and hence of FP. ●

7. The free skeleton of a sheet.

Although the analysis to now is sufficient to solve the motion-planning problem, it has one defect which is against the spirit of the Voronoi diagram approach: the skeleton includes the nonfree boundary of sheets, and therefore motions along the skeleton need to be perturbed to lie within FP. However, it is not necessary that this is so: for the remainder of this paper we shall analyze the role played by the semifree boundary and conclude that the free part of every skeleton is also connected. To avoid the cacophony 'nonfree semifree', in this section we use the term 'semifree' exclusively for placements in $SFP \sim FP$.

Definition. The *conjugate* of a placement $Z = (x,y,\theta)$ is the placement $(x+d\cos(\theta), y+d\sin(\theta), -\theta)$. In other words, it is the unique other placement Z' such that $B(Z') = B(Z)$.

Definition. Let Z be a virtual semifree placement in $VM(s_1,s_2)$. Let q and p be the points where the ladder touches (the extensions of) s_1 and s_2 respectively. Then the placement is *positive* if either (a) $p \neq q$ and q lies between p and Q on the ladder or (b) $p = q = P$. Otherwise the placement is *negative*.

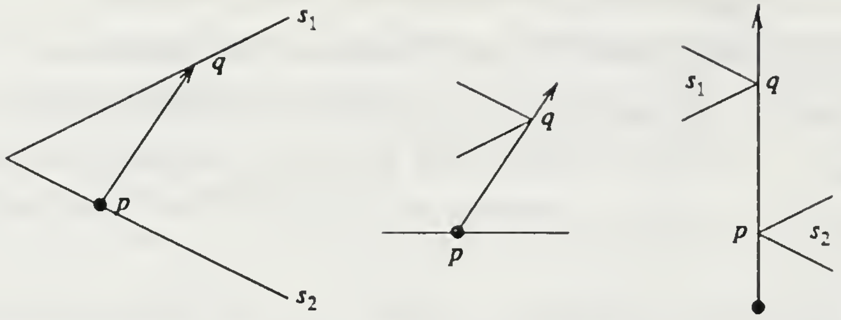

Figure 18. Positive semi-free placements

Note that this classification depends on the ordering of s_1 and s_2 – we follow the usual convention that in the wall-corner case s_1 is the corner. Note that case (b) is possible only when the ladder touches the apex of the wedge in the wall-wall case. It is clear that a placement is positive if and only if its conjugate is negative.

The basic fact which we shall establish is simply this: that in any sheet K, the intersection of the skeleton with FP is connected. This intersection we define as the *free skeleton* of a sheet. It is not easy to prove this fact in a general way, since the details are somewhat different for each of the three kinds of sheet; it is simplest to verify it case by case. Most of the work will use the fact that in any sheet both the set of positive semifree placements and the set of negative semifree placements is connected. Therefore we shall divide the analysis into three parts, one for each kind of sheet, and in each part we shall prove (i) that the set of positive semifree placements (and hence by symmetry the set of negative semifree placements) is connected, and then (ii) that the semifree skeleton is connected.

Lemma 7.1. If a sheet K contains a placement Z and its conjugate, then $\Theta(K)$ is the whole circle.

Proof. Let Z_t be a path in K from $Z = Z_0$ to its conjugate Z_1. Therefore $[\theta_0,\theta_1] \subseteq \Theta(K)$ where $\theta_i = \theta(Z_i)$. Then we may take the conjugate of every placement along the path and obtain a path from Z_1 to Z_0. This path must also lie in K since the clearance conditions are invariant with respect to conjugation. Hence, $[\theta_1,\theta_0] \subseteq \Theta(K)$. Q.E.D.

We shall see later that except possibly in the wall-corner case the sets of positive and negative semifree placements in a sheet (assuming they both exist) are mutually conjugate, essentially because the clearance function is invariant under conjugation.

Lemma 7.2. In any corner-corner sheet K, the set of positive semifree placements is a connected set of broadside placements. All semifree placements are in the same end of K.

Proof. The connectivity property is trivial, since the geometry is very simple. The second part is also very simple; note that it is possible that the set of semifree placements be common to both ends of the sheet (see below). ●

Lemma 7.3. The free skeleton of any corner-corner sheet K is connected.

Proof. We may assume that K has broadside placements. The proof is divided into three cases following the classification of the corners and its incident walls preceding Lemma 6.4 (see figure 12). These cases essentially classify the sheet based on the way in which the walls incident to the two corners extend from the y-axis.

(a) In this case all walls extend westward from the y-axis, and it is clear that the semifree placements are all in the west end of the sheet. (a.1) If there are two components of semifree placements then all orientations are feasible and the two ends are each homeomorphic to a circle; it follows that the intersection of FP with the west end of the sheet has two components corresponding to the ranges of angles $(-\pi/2,\pi/2)$ and $(\pi/2,-\pi/2)$. Each of these components intersects one of the two core streamlines which intersects the eastern end of the sheet. But the eastern end is connected and free. (a.2) If there is only one component of semifree placements, all positive placements, say, then again the western end may have up to two components of free placements corresponding to ranges of orientations $[\theta_{min},\pi/2)$ and $(\pi/2,\theta_{max}]$. Each of these components is connected to the eastern end of K by an extremal streamline. (Here the designated angles represent the bounds on $\Theta(K)$; the analysis holds whether or not either bound equals $\pi/2$.) If $\pi/2$ is a bound for the orientations then the extremal y-streamline represents the intersection of the east and west ends in an isolated nonfree placement. Otherwise both extremal streamlines consist only of free placements and connect the east end of K with the two components of free placements in the west end of K.

(b) Clearly, in this case, $\theta_{min} = \pi/2$, and $\theta_{max} > \pi/2$ (this follows from the principle of general position). Note that in this case the semifree placements are common to both ends of the sheet. Thus, the free part of the west end of the sheet may be represented as $\bigcup\{WE(\theta) : \theta \in (\pi/2,\theta_{max}]\}$ with a corresponding formula representing the free part of the east end of the sheet. Both these sets are path-connected by Lemma 6.6, and they are connected by the the θ_{max}-streamline which is free.

(c) This case is essentially the same as (b): as observed in Section 6, $\pi/2$ is an upper bound for $\Theta(K)$ (here there is exactly one semifree placement in the sheet). This concludes the analysis of the corner-corner case. Q.E.D.

Lemma 7.4. The set of positive semifree placements in a wall-wall sheet K is connected.

Proof. First note that the set of positive virtual semifree placements in $VM(s_1,s_2)$ is continuously parametrized by orientation θ in the range $[-\phi,\pi-\phi)$ where the angle ϕ defines the boundary of the wedge J defined in Section 3 for the standard parametrization. Note in the wall-wall case every nonfree placement is westernmost in its streamline. Let I be the set of orientations θ such that $W(\theta)$ is nonfree: we want to show that I is connected. Suppose not, so there exists two angles $\theta_0 < \theta_1$ in I such that $(\theta_0,\theta_1) \cap I = \varnothing$. First note that in order to reach a contradiction it is enough to show that there is some subinterval (ϕ_0,ϕ_1) of (θ_0,θ_1) such that both it and $(-\phi_1,-\phi_0)$ fail to intersect $\Theta(K)$, since any path from within the θ_0 streamline to the θ_1-streamline would contain placements in one of those intervals. Here we shall examine the orientations close to θ_0, analyzing why θ_0 is an upper bound for an interval of I, – choosing this angle for ϕ_0 – and deduce that these conditions are locally invariant under conjugacy. Let Z_0 be the semifree placement at orientation θ_0. There are several possible reasons that Z_0 should be at a locally maximal orientation:

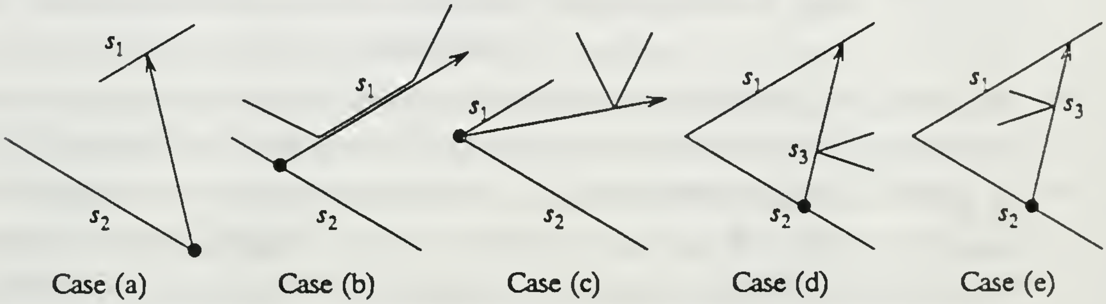

Figure 19. Illustrating the cases in Lemma 7.4

(a) P is at an eastern endpoint of s_2. Clearly, for all placements Z in $VM(s_1,s_2)$ such that $\theta(Z)$ exceeds θ_0 but is close to it, the virtual Z-racetrack does not meet s_2 and hence Z cannot be in K. The same considerations apply when Q touches a western endpoint of s_1.

(b) P is at the apex of the wedge J, $\theta_0 = \phi$ (the slope of s_1), and the ladder is longer than the distance from the apex of J to the eastern endpoint of s_1. Then for a placement Z chosen as in (a) the virtual racetrack fails to touch s_1 (or its endpoint).

(c) P is at the apex of the wedge, and the ladder touches a third object s_3. Clearly there exists an interval (θ_0,ϕ_1) of orientations for which for any placement $(0,0,\theta)$ with orientation in this interval, the ladder penetrates the interior of Ω^c. Thus if Z is any placement in $VM(s_1,s_2)$ with orientation θ then $\Gamma(Z)$ would contain some point p interior to Ω^c but not in $\Gamma((0,0,\theta_0))$, so by Lemma 5.4 such a placement cannot be in K.

(*d*) P touches s_2, Q touches s_1, and the ladder touches another object s_3, which must be a corner (general position). In this case we can prove directly that K does not admit any positive semifree placement whose orientation exceeds θ_0. We have two subcases. (*d*.1) Suppose that the walls incident to s_3 extend eastward from $B((0,0,\theta_0))$. Then there is an interval (θ_0,ϕ_1) with the property that for every virtual *semifree* placement Z in $VM(s_1,s_2)$ whose orientation is in this range, there exists some point p (depending on $\theta(Z)$) in $\Gamma(Z)$ but not in $\Gamma((0,0,\theta_0))$; hence by Lemma 5.6 no such placement can be in K. (*d*.2) If the walls extend westward from $B((0,0,\theta_0))$ then we can reach the same conclusion by considering the sets $\Sigma(Z)$.

This concludes the proof. Q.E.D.

Remarks. In case (*d*) of the above lemma the treatment was unlike that of other cases. In fact, the first part of case (*d*) could have been treated like the other cases: it would mean essentially that θ_0 is an upper bound for $\Theta(K)$. However, this is not true in the second part of case (*d*), since we can only deduce that θ_0 is an upper bound for I; it is still possible, for instance, that all orientations be admissible in the sheet. Compare with case (*e*) in Lemma 7.7. It is clear that the semifree placements always lie in the west end of the wall-wall sheet K. Note also that if case (*b*) or (*c*) applies then the semifree placements actually form the west end of K. For suppose the contrary, so there is some free placement Z in the west end of K. This can happen in only one way (since the walls s_1 and s_2 meet): there exist some further obstacle points in $\Gamma(Z)$. Therefore (by the 'thick walls' principle) there exist some points interior to Ω^c in $\Gamma(Z)$. This certainly cannot happen for any semifree placement satisfying (*b*) or (*c*), so Lemma 5.4 yields a contradiction.

Lemma 7.5. The free skeleton of a wall-wall sheet K is connected.

Proof. If the set of semifree placements has two components, then these components are mutually conjugate and therefore all orientations are feasible in K. Thus deleting the semifree placements from the west end yields two components (the free part of the west end also has two mutually conjugate components). By the remark following Lemma 7.4, each of these components intersects a core streamline which serves to connect them to the connected east end of K. This disposes of the case where the set of semifree placements has two components, so we may now assume that not all orientations are admissible in K and the set of semifree placements has exactly one component. Since we already know that the west end of K is homeomorphic to a line-interval we deduce that removing the set of semifree placements leaves two intervals each linked to the east end along an extremal streamline. Q.E.D.

It remains to consider the case of a wall-corner sheet. This is somewhat different from the other two cases in that it is possible for the free boundary of such a sheet to be disconnected without all orientations being admissible; in other words, it is possible for both extremal streamlines to terminate in semifree placements at opposite ends of the sheet. To see how this is so, let us begin the analysis by considering the set of positive virtual semifree placements in $VM(s_1,s_2)$ in the wall corner case. Such placements Z_θ exist for all orientations θ in the range $[\theta_{LB},\theta_{UB}]$, where the bounding angles are the positive solutions of the trigonometric relation

$$d\sin\theta = h$$

(Recall that h is the distance from s_1 to the extension of s_2.) The mapping $\theta \to Z_\theta$ is continuous and is a local parametrization except at the bounding angles. As θ varies over this range the end Q of the ladder traces out a quartic curve, a conchoid of Nicomedes [SS1], which has a cusp at s_1. The tangents to this cusp make angles θ_{LB} and θ_{UB} respectively. Suppose that the interior of this conchoid intersects Ω^c; it will follow that the ladder would pass out of *SFP* in tracing this curve. Let θ_1 be the first angle at which this begins to happen (it is possible that θ_1 equal θ_{LB}, i.e., when the line through the ladder penetrates the interior of Ω^c at s_1), and let θ_2 be the last angle at which this happens.

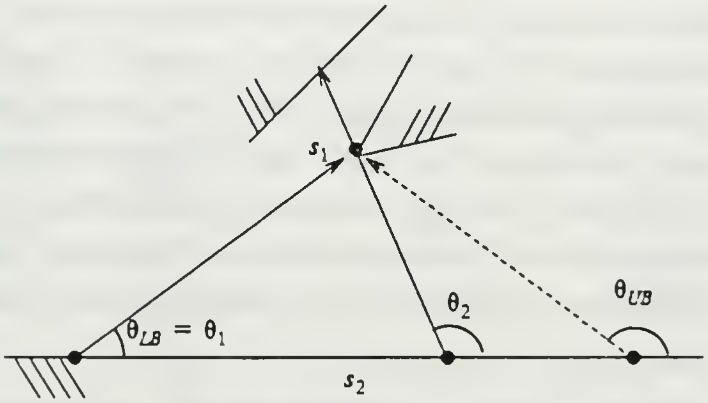

Figure 20. Illustrating the bounding angles θ_1 and θ_2

Lemma 7.6. Following the above discussion for the wall-corner sheet K, if θ_1 and θ_2 are both defined, and $\Theta(K)\cap[-\theta_2,\theta_1] \neq \varnothing$, then $\Theta(K) \subseteq [-\theta_2,\theta_1]$.

Proof. It is not difficult to show that there exist angles ϕ_1 and ϕ_2 close to θ_1 and θ_2 respectively such that for any placement Z such that $\theta(Z)$ or $-\theta(Z)$ is in the range $(\theta_1,\phi_1) \cup (\phi_2,\theta_2)$ and Z is in $VM(s_1,s_2)$ then Z cannot be in any (s_1,s_2) sheet. To show this one merely relies on the 'thick walls' principle and the growth properties of the sets $\Delta(Z)$ as used in demonstrating connectivity of streamlines.

However, any continuous path from a placement Z in $VM(s_1,s_2)$ whose orientation is in $[-\theta_2,\theta_1]$ to a placement in $VM(s_1,s_2)$ whose orientation is outside that range would have to contain a placement whose orientation lay in one of these four 'forbidden' intervals. Q.E.D.

Lemma 7.7. The set of positive semifree placements in a wall-corner sheet K is connected.

Proof. As in the wall-wall case, let I be the set of angles defining semifree placements in K and suppose that θ_0 and $\theta_1 \in I$ but $(\theta_0,\theta_1) \cap I = \emptyset$. Again it will be enough to show that there is always some ϕ_1 in (θ_0,θ_1) such that for no placement Z in $VM(s_1,s_2)$ such that Z or its conjugate has orientation in (θ_0,ϕ_1) is Z in K. The following situations are possible.

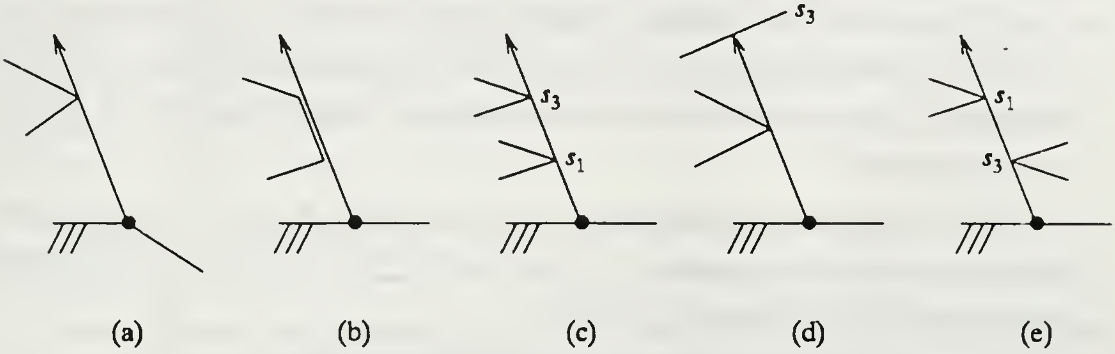

Figure 21. Illustrating the cases for Lemma 7.7

(*a*) P touches the eastern endpoint of s_2. Clearly, for all θ greater than θ_0 and sufficiently close to it, if Z is a placement at orientation θ in $VM(s_1,s_2)$ then either (i) the ladder is west of s_1 in which case the virtual racetrack must intersect the interior of Ω^c or (ii) the ladder is east of or touches s_1 and the virtual racetrack fails to touch s_2.

(*b*) the ladder lies along a wall incident to s_1; (*c*) the ladder touches another corner s_3 where the point of contact is above the point where it touches s_1; or (*d*) Q touches another wall s_3. All of these cases have been covered in Lemma 7.6.

(*e*) The ladder touches another object s_3 between P and the point of contact with s_1. By the principle of general position, s_3 must be a corner. Both walls incident to s_3 must extend eastward from that corner. In this case we can show directly that there can be no positive *semifree* placement in the sheet with orientation greater than θ_0. Let Z_i represent the positive semifree placements with orientations θ_i. Clearly $\Gamma(Z_1)$ contains some point p interior to Ω^c which is not in $\Gamma(Z_0)$. This directly contradicts Lemma 5.4.

This concludes the proof. Q.E.D.

Lemma 7.8. The free skeleton of any wall-corner sheet K is connected.

Proof. If either (*i*) all orientations are admissible or (*ii*) the set of semifree placements is connected then the arguments are the same as for the other two kinds of sheet; hence let us only consider the case where the set of semifree placements is disconnected and not all orientations are admissible. Let θ_{min} and θ_{max} be the bounds on $\Theta(K)$. We know that these angles have opposite signs since the sheet has both positive and negative semifree placements. Claim that both θ_{min} and θ_{max} admit semifree placements; it is enough to discuss this for θ_{max}. The only case in which this might fail is when the angle θ_0 in the above Lemma resulted from case (*e*). However, $\theta_{max}-\theta_{min} < \pi$. This implies that the region in Ω bounded by the x-axis and two lines through s_1 at these two angles contains s_3 and points near s_3 interior to Ω^c. This is impossible, since such points would be in $\Gamma(Z_1)\sim\Gamma(Z_0)$ where Z_0 is the semifree placement at orientation θ_0 and Z_1 is any negative semifree placement.

Thus deleting the semifree placements from an end of K leaves two components: one extremal streamline which intersects the other end in a free placement, and one connected set which intersects the core streamline. Q.E.D.

Summarizing the results of this section is the following theorem.

Theorem 7.9. In any sheet K, the free skeleton of K, i.e., the intersection of FP with the union of ∂K and its core streamlines, is connected. ●

Define the *free skeleton* N of FP to be the union of the free skeletons of its sheets. The concluding result of this paper is that N defines the path-connectivity of FP.

Theorem 7.10. Every path-connected component of FP contains a unique path-connected component of N.

Proof. Let Z_0 and Z_1 be in the same path connected component of FP. By the retraction theorem, this is equivalent to saying that $Im(Z_0)$ and $Im(Z_1)$ are in the same path-connected component of $Vor_d(\Omega)$. Let Π be a path from $Im(Z_0)$ to $Im(Z_1)$ in $Vor_d(\Omega)$ and let H be the collection of sheets visited along this path. For any placement Z in the path not interior to any sheet, Z must be in the free skeletons of all adjacent sheets. Hence the union of the free skeletons of all sheets in H is connected and therefore contained in the same component of N. The converse is trivial since $N\subseteq FP$. Q.E.D.

8. Remarks on constructing and using the free skeleton for motion-planning.

In this section the phases of the motion-planning algorithm are sketched out. The companion paper [OSY2] describes an efficient implementation of the preprocessing stage. This preprocessing amounts, of course, to constructing a combinatorial representation of the free skeleton N. Recall that N is a network consisting of *ends* and *core streamlines*. The ends of a sheet can be divided into subparts called *edges* where each is bounded by *Voronoi vertices* which represent placements common to four sheets, and *semifree vertices* (which are not, strictly speaking, in N). We shall call placements at the intersection of core streamlines with edges (respectively, with the semifree boundary of FP) *pseudo-vertices* (respectively, *semifree pseudo-vertices*). In general, edges may be associated with three objects s_1,s_2,s_3 and represent placements in $M(s_1,s_2,s_3)$. We retain $VM(s_1,s_2,s_3)$ to denote the virtual midline defined by them. Suppose that θ is an orientation which is not degenerate in the sense of Section 6. In general, when the midline is associated with separated objects one of which is a wall, there may be two placements Z in $VM(s_1,s_2,s_3)$ at orientation θ. There cannot be more than two since otherwise the same three objects would define three Voronoi vertices in the standard Voronoi diagram $Vor_0(\Omega_\theta)$ which is known to be false by considerations of planarity of the standard Voronoi diagram. It is significant, however, that not both placements can be in $M(s_1,s_2,s_3)$.

Lemma 8.1. For any orientation θ, the set of placements Z at orientation θ in $M(s_1,s_2,s_3)$ is connected.

Proof. Consider two placements Z_1 and Z_2 in $VM(s_1,s_2)$ in which the (virtual) racetrack touches a third separated objects s_3 at orientation θ, and first suppose that the placement is not degenerate (i.e., broadside with respect to s_3 and one of the other objects). Let L be the common tangent to both these racetracks. Clearly, both racetracks must touch s_3 at opposite sides of L, and hence, by the thick walls principle, they cannot both be in $M(s_1,s_2,s_3)$. Similar considerations apply in the case where the placements are broadside with respect to s_3 and, say, s_1: in placements Z such that Z is in $M(s_1,s_2,s_3)$ and broadside with respect to the two named objects, the Z-racetrack must always be on the same side of the line L joining these two corners, and it only requires a short case-analysis to verify that in this context too the set of such placements is connected. Q.E.D.

The next observation is that any nonempty core streamlines may be associated with a unique edge of the standard Voronoi diagram $Vor_0(\Omega)$.

Lemma 8.2. Any core streamline in an (s_1,s_2)-sheet may be associated with a unique edge corresponding to the same two (separated) objects in $Vor_0(\Omega)$. Furthermore, with every edge e of $Vor_0(\Omega)$ either zero or two core streamlines are associated where the two core streamlines need not be in the same sheet.

Proof. Let X denote the given core streamline. Because in any placement Z on this streamline the ladder is horizontal with respect to the standard coordinate system for $VM(s_1,s_2)$, it follows (by a short case-analysis) that there is a unique point $p(Z)$ on $B(Z)$ simulaneously closest to s_1 and s_2, and $p(Z)$ varies continuously with Z. Furthermore, the disc $D(Z)$ of radius $Clearance(Z)$ about $p(Z)$ touches s_1 and s_2 and is contained in the racetrack $R(Z)$. Hence, for all placements Z in the streamline, $p(Z)$ is confined to the same edge e of $Vor_0(\Omega)$. This is the unique edge with which the core streamline may be associated. Next suppose that X represents the virtual streamline extending the core streamline in $VM(s_1,s_2)$. If both objects s_i are corners, then by the analysis in Lemma 5.2 we know that there is only one (s_1,s_2) segment of this streamline intersecting $M(s_1,s_2)$, and there is nothing more to prove (this is analogous to the fact that two corners can determine at most one edge of the standard Voronoi diagram).

Next suppose that both objects are walls. Suppose that the virtual streamline intersects more than one (s_1,s_2)-sheet. We can thus suppose that there exist placements Z_1 and Z_2 in the virtual streamline, where Z_1 is the eastern endpoint of one sheet and Z_2 is the western endpoint of the 'next' sheet, indicating that no intermediate placements are in the closure of $M(s_1,s_2)$. From the analysis in Lemma 5.5, the only way this can happen is for there to exist points p interior to Ω^c and in $\Sigma(Z_1)$ but not in $\Sigma(Z_2)$. Observe, however, that $\Sigma(Z)$ contains the disc $D(Z)$ introduced above. It follows by simple 'sweeping' arguments that there exists a placement Z_3 between these two placements (actually, at Z_2), such that $D(Z_3)$ touches a third object. Hence Z_1 and Z_2 are associated with different edges of $Vor_0(\Omega)$.

Finally, let us consider the wall-corner case. With Z_1 and Z_2 as in the wall-wall case, we apply the same arguments to deduce the existence of a placement Z_3 between them where $P(Z_3)$ is simulaneously closest to three objects, and again Z_1 and Z_2 cannot ba associated with the same edge of $Vor_0(\Omega)$. Q.E.D.

Finally we can briefly review the preprocessing and path-planning aspects of the algorithm. We may suppose that we have at our disposal (i) the standard Voronoi diagram $Vor_0(\Omega)$, and (ii) all the Voronoi vertices and semifree vertices. (Efficient 'collection' of these vertices can be implemented in time $O(n^2\log(n)\log^*(n))$, with some effort [OSY2].) Each of these vertices and semifree vertices can be associated with the (one or four) sets of the form $M(s_1,s_2,s_3)$ corresponding to the incident Voronoi edges. Note that the given midlines may, by Lemma 8.1, be parametrized by a parameter t, (usually based on orientation, but with provision made to parametrize degenerate placements). Thus these sets may be assumed sorted in ascending (or cyclic) order based on such a parametrization. At this point any pseudo-vertices may be introduced by simply examining the range of orientations in the edges of $M(s_1,s_2,s_3)$. Finally, vertices at opposite ends of the same core streamline may be identified by referring to the standard Voronoi diagram (Lemma 8.2). Thus we can suppose that the network N has been assembled.

Thus, given two free placements Z_1 and Z_2, to find a path between them (if one exists) one can proceed as follows.

(i) Compute $Im(Z_i) = Y_i$, $(i = 1,2)$. This can be done by naive methods in time $O(n)$.

(ii) Locate the two endpoints of the streamlines containing the respective placements Y_i. Again, this can be done naively in time $O(n)$. At most one of these endpoints in each streamline can be nonfree, so by this means one can locate placements in the (topological) skeleton of the sheet. Thus one can determine by geometric means placements in one or two midlines of the form $M(s_1,s_2,s_3)$.

(iii) Finally, to begin the search of the combinatorial network N, simply find (again, naively in time $O(n^2\log^* n)$) the edges in the (sorted) representations of the associated midlines. From their the network can be searched as a combinatorial graph. Note that the network has size $O(n^2\log^* n)$ [OSY] so this search will dominate the runtime.

REFERENCES

[HJW] J. E. Hopcroft, D. A. Joseph, and S. H. Whitesides, "On the movement of robot arms in 2-dimensional bounded regions," *23rd FOCS*, pp. 280-289, 1982.

[Kir] D. Kirkpatrick, "Efficient Computation of Continuous Skeletons," *20th FOCS*, pp. 18-27, 1979.

[LW] T. Lozano-Perez and M. Wesley, "An algorithm for planning collision-free paths among polyhedral obstacles," *CACM*, vol. 22, pp. 560-570, 1979.

[OY] C. Ó 'Dúnlaing and C. K. Yap, "A retraction method for planning the motion of a disc," *J. Algorithms (to appear).*

[OSY2] C. Ó 'Dúnlaing, M. Sharir, and C. K. Yap, "Generalized Voronoi Diagrams for a Ladder: II. Efficient construction of the diagram," Courant Institute and Robotics Lab. Report, September 1984 .

[OSY1] C. Ó 'Dúnlaing, M. Sharir, and C. K. Yap, "Retraction: a new approach to motion-planning," *15th STOC*, pp. 207-220, 1983.

[Row] P. F. Rowat, "Representing spatial experience and solving spatial problems in a simulated robot environment," *Ph.D. Thesis*, University of British Columbia, 1979.

[SS1] J. T. Schwartz and M. Sharir, "On the piano movers' problem: I. The special case of a rigid polygonal body moving amidst polygonal barriers," *Comm. Pure Appl. Math.*, vol. XXXVI, pp. 345-398, 1983.

[SS2] J. T. Schwartz and M. Sharir, "On the piano movers' problem: II. General Techniques for computing topological properties of real algebraic manifolds," *Advances in Applied Math.*, vol. 4, pp. 298-351, 1983.

APPENDIX: Growth properties of the sets $\Sigma(Z)$, $\Gamma(Z)$, and $\Delta(Z)$

This appendix proves lemma 3.3 concerning certain monotonicity properties of the sets $\Sigma(Z)$, $\Gamma(Z)$ and $\Delta(Z)$. For the reader's convenience, we summarize the lemma here:

Lemma. As Z moves along a virtual θ-streamline parametrized by the x-coordinate, the set $\Gamma(Z)$ is expanding; the set $\Sigma(Z)$ is contracting; and the set $\Delta(Z)$ is contracting for Z west of a certain point x_0 and expanding east of x_0. In each case, the expansion and contraction are strict except for the following set Y of points, defined for each type of sheet:

- [CC] Y is the set of points on the y-axis, together with the points on any lines through those sides (if any) of the Z-racetrack that touch s_1 or s_2.
- [WW] Y is the set of points on the extended walls.
- [WC] Y is the set of points on the horizontal lines $y = 0$ and $y = h$, together with the points on any lines through the side (if any) of the Z-racetrack that touches s_1.

This result is proved in three parts, as lemmas A.4, A.5, and A.6 below. First we note three simple geometrical lemmas. For any x, let p_x denote the point $(x, 0)$ on the x-axis.

Lemma A.1. Let q_1 and q_2 be antipodal points on the y-axis. For each x define D_x to be the disc centered at p_x and passing through q_1 and q_2, and define γ_x (respectively, σ_x) to be the intersection of D_x with the half-plane east (resp. west) of the y-axis. Then $\gamma(x)$ (resp. $\sigma(x)$) expands (resp. contracts) as x ranges over the whole line. Moreover the expansion (resp. contraction) is strict except for the points on the y-axis.

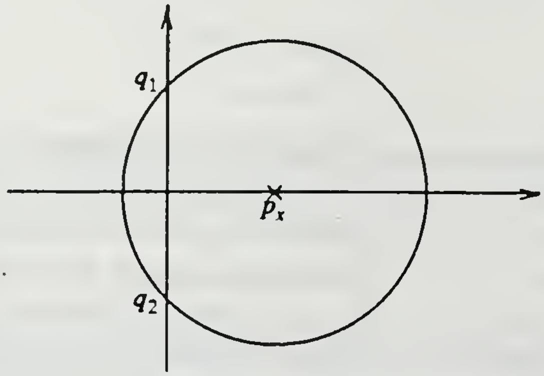

Figure A.1

We omit the simple proof since it is similar to the next.

Lemma A.2. Let L_1 and L_2 be two rays extending from the origin at angles $\pm\phi$, where $0 < \phi < \pi/2$. As usual let J denote the closed convex wedge enclosed by the rays. For every $x \geq 0$, let D_x be the closed disc centered at p_x and touching both rays; let γ_x be the union of D_x with the set of all points in the wedge J and lying west of D_x, and let σ_x be the union of D_x with the set of all points in the wedge east of D_x. Then γ_x expands and σ_x contracts for $x \geq 0$. This expansion (resp. contraction) is strict except on the rays L_1 and L_2.

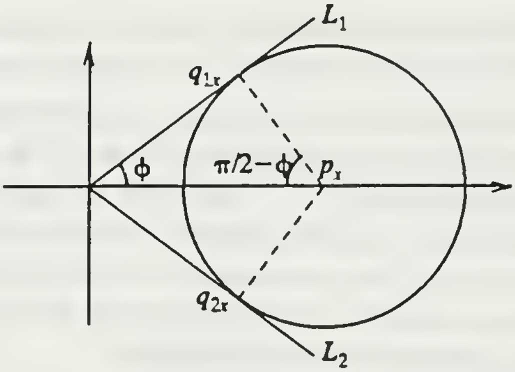

Figure A.2

Proof. For any $x > 0$ let q_{ix} $(i = 1,2)$ be the point where D_x touches L_i. If a, b, c are any non-collinear points, we let $\angle abc$ denote the angle at b subtended by a, c, measured counter-clockwise from a to c. A point p $(\neq q_{1x}, q_{2x})$ in the wedge J is in γ_x if and only if

$$\angle q_{1x} p q_{2x} \geq \pi/2-\phi.$$

Moreover this inequality is strict if and only if p does not lie on the curved portion of the boundary of γ_x. To show that γ_x is expanding, suppose $u < x$. Since the angle $\angle q_{1x} p q_{2x}$ increases strictly with increasing x, we see that p is in γ_u implies p is in γ_x. Moreover, p is in the interior of γ_x unless p lies on the rays. A similar analysis holds for σ_x based on the inequality

$$\angle q_{2x} p q_{1x} \geq \pi/2+\phi.$$

Q.E.D.

The constructions involved in the next lemma are as follows: let q_1 be a point $(0, h)$ on the positive y-axis and let D_x represent the closed disc touching q_1 and tangential to the x-axis at p_x (so the center of D_x is on the parabola with q_1 as focus and the x-axis as directrix). Let A be the infinite strip $0 \leq y \leq h$. The line segment q_1p_x divides A into an eastern and a western component denoted E_x and W_x. Let γ_x (resp. σ_x) denote $D_x \cap E_x$ (resp. $D_x \cap W_x$). Also let δ_x be the intersection of D_x with the half-space $y \geq h$. The next lemma is a variation on the expanding and contracting set idea.

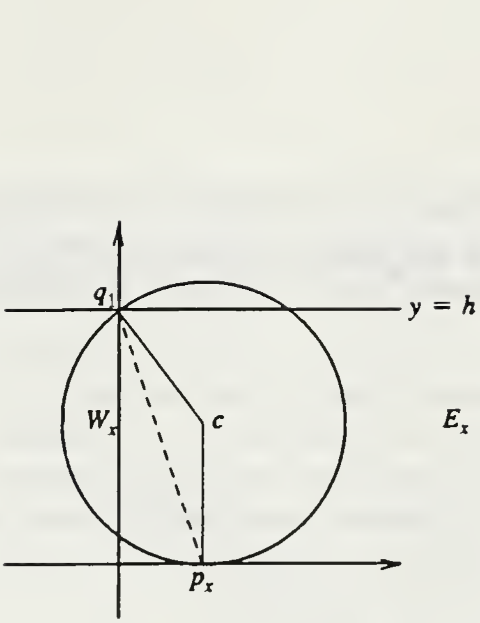

Figure A.3 γ_x, σ_x and δ_x.

Lemma A.3. Suppose $u < x$. Then

(i) $p \in \gamma_u \cap E_x$ implies $p \in \gamma_x$. Moreover, if p does not lie on the line $y = h$ then p is in the interior of γ_x.

(ii) $p \in \sigma_x \cap W_u$ implies $p \in \sigma_u$. Moreover, if p does not lie on the line $y = h$ then p is in the interior of σ_u.

(iii) The set δ_x is contracting for $x \leq 0$ and expanding for $x \geq 0$. Moreover, the expansion and contraction is strict except for points on the line $y = h$.

Proof. For each x let α_x be the angle $\angle q_1 c p_x$ where c is the center of D_x. Clearly α_x decreases monotonically with increasing x. Note that for any point $p \neq q_1$ in E_x, p lies in γ_u if and only if

$$\angle q_1 p p_u \geq \frac{\alpha_u}{2}.$$

Moreover, if p is not on the line $y = h$ then the inequality is strict iff p is in the interior of γ_u. To see (i), note that $p \in \gamma_u \cap E_x$ implies

$$\angle q_1 p p_x > \angle q_1 p p_u \geq \frac{\alpha_u}{2} > \frac{\alpha_x}{2}.$$

This clearly implies that p is in γ_x; moreover, p is in the interior of γ_x if p is not in the line $y = h$. The proof for σ_x is symmetric.

To show (iii) for δ_x, we note that for $x \geq 0$ (resp. $x \leq 0$), the arguments are as for γ_x (resp. σ_x). **Q.E.D.**

In the following, the ladder moves along a θ-streamline V_θ (or y-streamline V_y in the case of broadside placements) parametrized as Z_x for all x (or for $x \geq c$ in the $\boxed{\mathrm{WW}}$ case, c depending on θ). We adopt the notation Γ_x, Σ_x, and Δ_x for $\Gamma(Z_x)$, $\Sigma(Z_x)$, and $\Delta(Z_x)$ respectively.

Lemma A.4. (WW) Let Z_x be the standard parametrization of a virtual streamline V_θ in the virtual midsurface of two walls s_1 and s_2. Then the set Γ_x is expanding, and Σ_x contracting, with x. Moreover, except for points on the extended walls, the expansion and contraction are strict.

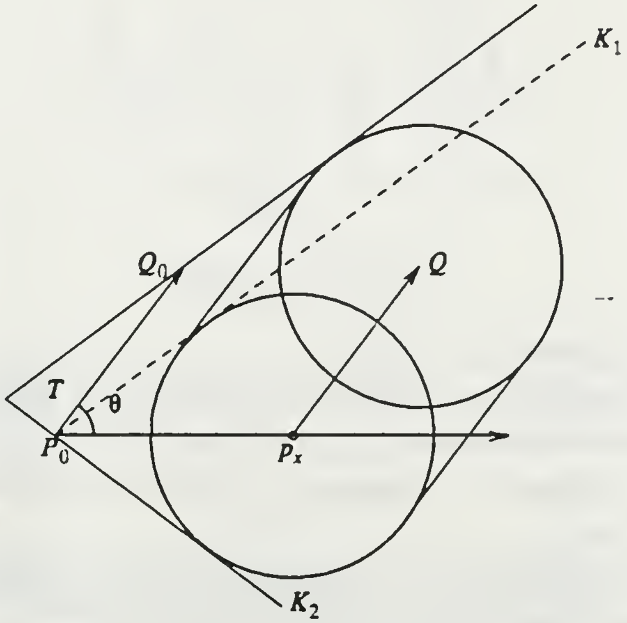

Figure 4. Wall-wall case.

Proof. We prove this for Γ_x only since the other case is similar. Without loss of generality assume $0 \leq \theta < \pi$. Note that if $\theta < \pi-\phi$ then both ends, $P_0 = P(Z_0)$ and $Q_0 = Q(Z_0)$, of the ladder at the western extreme of the streamline V_θ lie on the extended walls; let T denote the triangle ΔOP_0Q_0. Otherwise, if $\theta < \phi$ (resp. $\pi-\phi < \theta$) then P_0 (resp. Q_0) coincides with the origin O; T is defined to be the single point O in this case. Let K_1 and K_2 be the rays emanating from P_0 at angles $\pm\phi$ – so K_i are the displaced walls. So K_i is parallel to the wall s_i. With respect to the K_i's, the sets D_x, σ_x, and γ_x can be defined as in lemma A.2. Note that

$$\Gamma_x = T \cup (\gamma_x + B_\theta).$$

By lemma A.2, the set γ_x is expanding with x and is strict except on K_1 and K_2. This translates to the claim that Γ_x is expanding with x and that the expansion is strict except on the extended walls. **Q.E.D.**

Lemma A.5. (CC). Let Z_x be the standard parametrization of the virtual streamline V_θ (or V_y if the ladder is broadside) in the virtual midsurface of two corners s_1, s_2. Then the sets Γ_x and Σ_x respectively expand and contract with x. Moreover the expansion and contraction are strict except for the points on the line s_1s_2 and on any line through those sides of the racetrack that touch s_1 or s_2.

Proof. We shall consider only θ-streamlines, assuming that the ladder is not vertical; the case where the ladder is vertical is even simpler.

Consider the line-segments $s'_i = s_i - B_\theta$ $(i = 1, 2)$: let $t_i = s_i - (d\cos\theta, d\sin\theta)$ denote the other endpoint of s'_i. It is easily seen that the P-end of a virtual streamline V_θ traces out the standard Voronoi diagram of $s'_1 \cup s'_2$. See figure A.5. In general, this trace may be divided naturally into 5 sections: three line-segments connected by two parabolic segments. The central segment disappears if θ is such that two normals to the s'_i's (through the points s_2 and t_1) coincide, but this does not affect the generality of our analysis.

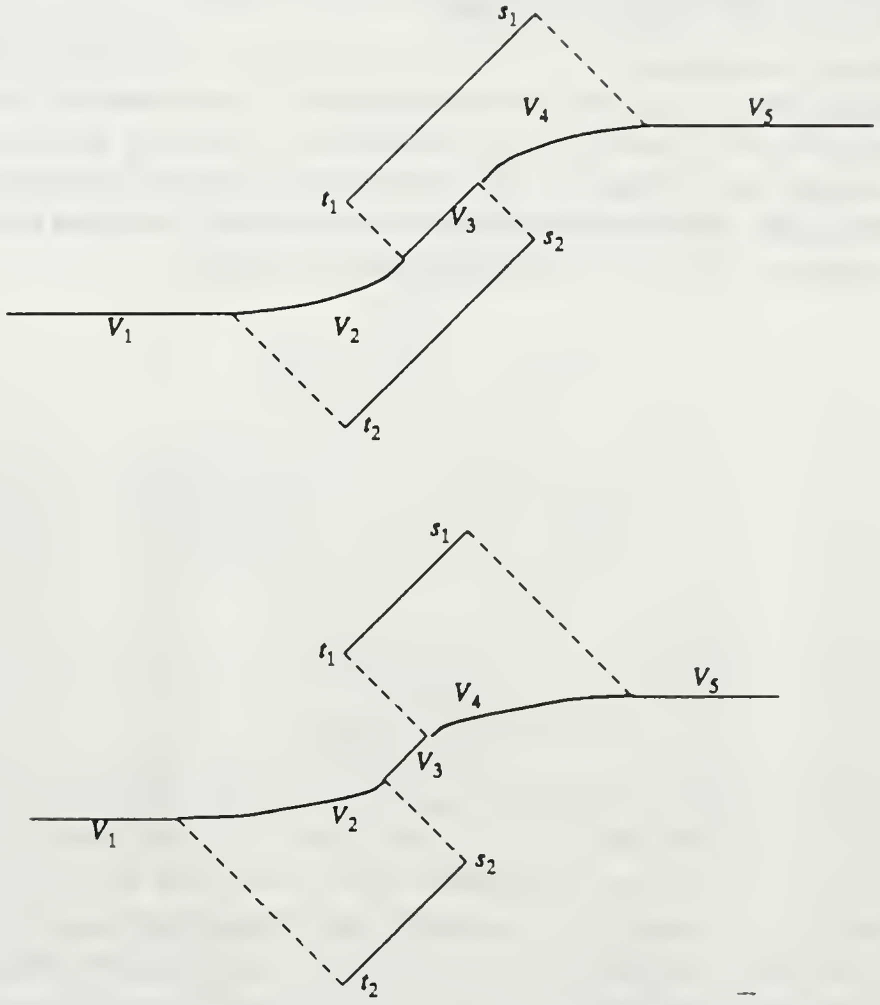

Figure A.5. Two instances of the segmentation of (the trace of) V_θ into five parts.

Let us divide the streamline into segments V_1 through V_5, corresponding to the segmentation of its trace. Thus let

$$-\infty = x_0 < x_1 < x_2 < x_3 < x_4 < x_5 = \infty$$

such that Z_τ lies in V_i iff $x_{i-1} \le x \le x_i$. In particular, V_1 (resp. V_5) is a ray that extends horizontally to the west (resp. east). Without loss of generality, $0 \le \theta < \pi/2$; the other cases are symmetric. Let H_W and H_E denote the half-planes to the west and east (respectively) of the line $s_1 s_2$. Considering each V_i in turn:

(V_1) V_1 lies along the bisector of the two points t_i $(i=1, 2)$. To analyze Γ_τ, define the sets

γ_x, σ_x and D_x with respect to the two points t_i (ie. regard t_i as antipodal points of the appropriate coordinate axis), as in lemma A.1. It is not hard to see that $\Gamma_x = H_E \cap (\gamma_x + B_\theta)$, and so lemma A.1 implies that Γ_x is expanding for $x \leq x_1$. Moreover, the expansion is strict except for points on the y-axis. The corresponding result for Σ_x follows by noting that $\Sigma_x = H_W \cap (\sigma_x + B_\theta)$.

(V_2) This corresponds to a parabolic segment generated by the point t_1 and the line segment s'_2. So for Z in V_2, s_1 touches the eastern bend and s_2 the southern side of the Z-racetrack. Let L_i ($i = 1, 2$) be the line through s_i with orientation θ and let p_t denote the point on L_2 whose x-coordinate is t. By choosing the coordinate system such that L_2 is the x-axis and s_1 corresponds to q_1 of lemma A.3, we can define the sets D_t, E_t, W_t, γ_t, σ_t, and δ_t. (For instance, D_t is the disc touching s_1 and tangent to L_2 at p_t.) Now observe

$$\Gamma_x = H_E \cap (\gamma_x + B_\theta)$$

To see that Γ_x is expanding, suppose $x < u$ and $p \in \Gamma_x$. If $p \in H_E \cap (W_u + B_\theta)$ then clearly $p \in \Gamma_u$. Otherwise $p \in (\gamma_x \cap E_u) + B_\theta$ and lemma A.3 implies that $p \in \gamma_u + B_\theta \subseteq \Gamma_u$. Similarly, Σ_x is contracting based on the observation

$$\Sigma_x = H_W \cap ((\sigma_x \cup \delta_x) + B_\theta),$$

and from lemma A.3. (in particular, δ_x is contracting for x in V_2). In each case, the expansion or contraction is strict except for points on the line L_2 or on the vertical line through $s_1 s_2$.

(V_3) This segment, if it exists, corresponds to the case where the corners s_i ($i = i, 2$) are either both at opposite bends of the racetrack or both at opposite sides of the racetrack. Since the geometry is very simple here, the reader is invited to supply the details. Finally, V_4 and V_5 are handled analogously to V_2 and V_1, respectively. **Q.E.D.**

Lemma A.6 (WC). Let Z_x be a parametrization of the streamline V_θ in a wall-corner virtual midsurface. The sets Γ_x and Σ_x respectively expand and contract for x ranging over the entire line. The set Δ_x contracts for $x \leq x_0$ and expands for $x \geq x_0$ where Z_{x_0} corresponds to the placement of smallest clearance on the streamline. Moreover, the expansion and contraction in each case are strict except at the horizontal lines through the origin and the corner s_1, and the side (if any) of the race-track that touches s_1.

Proof. We may assume that $0 < \theta \leq \pi/2$. Again, the P-end of the ladder in the streamline traces out the standard Voronoi diagram generated by the x-axis and the line segment $s'_1 = s_1 - B_\theta$. See figure A.6. It separates naturally into 5 closed line-segments V_i, parabolic for i=1, 3, 5, and straight for i=2, 4. If the streamline contains a semifree placement (ie. when $t_1 = s_1 - (d\cos\theta, d\sin\theta)$ lies below the x-axis) then the segment V_3 degenerates to a point; again, this does not affect the generality of our proof. For each x, let the disc D_x be centered at $P(Z_x)$ and have radius *Clearance*(Z_x). This disc is tangential to the x-axis and touches (not necessarily at a tangent) s'_1; let p_x and q_x respectively denote these two points of contact. Let A be the strip between the x-axis and the horizontal line through s_1. The strip is divided by the segment s_1p_x into an eastern and western part, E_x and W_x. Clearly E_x (resp. W_x) is contracting (resp. expanding). Considering each V_i in turn:

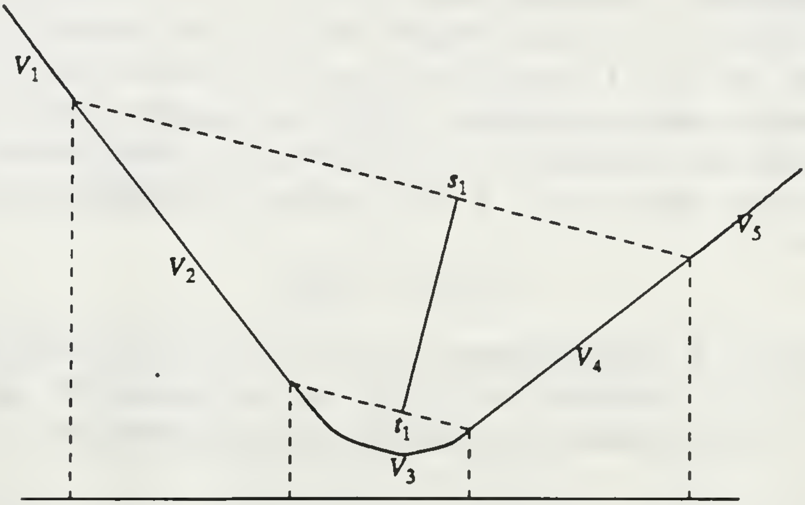

Figure A.6. The segmentation of the V_θ streamline
in the wall-corner case.

(V_5) First suppose that Z_x is in V_5 so the disc D_x touches the corner s_1. Note that Γ_x may be expressed as the union of W_x and $(D_x + B_\theta) \cap A$. Let $x < u$ and p be in Γ_x. If $p \in W_u \cup (B_\theta + (D_x \cap E_u))$ then lemma A.3 implies that $p \in \Gamma_u$. So we may suppose that $p \in B_\theta + (D_x \cap W_u)$. Then p has the form $q + r$ where $q \in D_x \cap W_u$, $r \in B_\theta$. Either $p \in W_u$ or else for some $0 < \alpha < 1$, $q + \alpha r$ lies on the line-segment s_1p_u. In the latter case $p = (q + \alpha r) + (1-\alpha)r$ lies in $(D_x \cap E_u) + B_\theta \subseteq \Gamma_u$. Thus we have shown that Γ_x is expanding. It is also easy to see that the expansion is strict except for points the set described by the lemma. The analysis for Σ_x is similar: Σ_x is the union of E_x and $(D_x + B_\theta) \cap A$, and our result follows as usual using lemma A.3. Finally, Δ_x is given by $\delta_x + B_\theta$ which (lemma A.3 shows) is expanding for $x \geq x_4$.

(V_4) We choose an appropriate coordinate system as in lemma A.2 and define the sets γ_x and σ_x. Then the desired growth properties easily follows from the observation that $\Gamma_x = W_x \cup (A \cap (\gamma_x + B_\theta))$, $\Sigma_x = E_x \cup (A \cap (\sigma_x + B_\theta))$, and Δ_x is the intersection of the half-plane $y \geq h$ with $\gamma_x + B_\theta$.

(V_3) The analysis for V_3 is similar to that for V_5 except that the disc D_x would touch $t_1 = s_1 - (d\cos\theta, d\sin\theta)$ rather than s_1. Also Δ_x is initially contracting then expanding in this range. Finally, the analysis for V_1 and V_2 can similarly be carried out. **Q.E.D.**

This book may be kept

FOURTEEN DAYS

A fine will be charged for each day the book is kept overtime.

GAYLORD 142 PRINTED IN U.S.A.

NEW YORK UNIVERSITY
COURANT INSTITUTE-LIBRARY
251 Mercer St. New York, N.Y. 10012

c.2

NYU Compsci TR-139
O'Dunlaing c.2

Generalized Voronoi

c.2

NYU Comp Sci TR-139
O'Dunlaing c.2

Generalized Voronoi
diagrams for moving

DATE DUE	BORROWER'S NAME

LIBRARY
N.Y.U. Courant Institute of
Mathematical Sciences
251 Mercer St.
New York, N. Y. 10012